United States Air Force
Scientific Advisory Board

Report on

Database Migration for Command and Control

SAB-TR-01-03

November, 2002

Cleared for Open Publication, September 2002

REPORT DOCUMENTATION PAGE	Form Approved OMB No. 0704-0188

Public reporting burder for this collection of information is estibated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burder to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (FROM - TO)
01-06-2002	Technical	xx-01-2001 to xx-11-2001

4. TITLE AND SUBTITLE Database Migration for Command and Control Unclassified	5a. CONTRACT NUMBER 5b. GRANT NUMBER 5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S) Hendler, James ; Author Thuraisingham, Bhanvani ; Author Nelson, Eric ; Author Wiederhold, Gio ; Author Saunders, Scott ; Author	5d. PROJECT NUMBER 5e. TASK NUMBER 5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME AND ADDRESS USAF Scientific Advisory Board Room 5D982 1180 AF Pentagon Washington, DC20330-1180	8. PERFORMING ORGANIZATION REPORT NUMBER SAB-TR-01-03
9. SPONSORING/MONITORING AGENCY NAME AND ADDRESS SAF/OS AF/CC Pentagon Washington, DC20330-1670	10. SPONSOR/MONITOR'S ACRONYM(S) SAF/OS; AF/CC 11. SPONSOR/MONITOR'S REPORT NUMBER(S) SAB-TR-01-03

12. DISTRIBUTION/AVAILABILITY STATEMENT
APUBLIC RELEASE
,

13. SUPPLEMENTARY NOTES

14. ABSTRACT
A report on the technologies, techniques and procedures required to manage effectively the transition of Air Force databases from one generation of information technology to the next.

15. SUBJECT TERMS
Science and Engineering Planning, Science and Technology, Leveraging Partnerships, Database Technology, Development Planning, Database Migration

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19. NAME OF RESPONSIBLE PERSON
			Same as Report (SAR)	156	Robert, Ripperger robert.ripperger@pentagon.af.mil
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER International Area Code 011 Area Code Telephone Number 703692-5097 DSN 222-5097

(This Page Intentionally Left Blank)

This report is a product of the United States Air Force Scientific Advisory Board Committee on Database Migration for Command and Control. Statements, opinions, recommendations, and conclusions contained in this report are those of the committee and do not necessarily represent the official position of the U.S. Air Force or the Department of Defense.

(This Page Intentionally Left Blank)

United States Air Force
Scientific Advisory Board

Report on

Database Migration for Command and Control

SAB-TR-01-03

November 2002

Cleared for Open Publication, September 2002

(This Page Intentionally Left Blank)

Table of Contents

(This Page Intentionally Left Blank)

List of Figures

List of Tables

(This Page Intentionally Left Blank)

Chapter 1

Introduction

Slide 1: Our Team

Our Team

Chair: Dr. James Hendler
Deputy Chair: Dr. Bhavani Thuraisingham
Civilian Advisor: Mr John Gilligan

Operations
Maj Gen (ret) Eric Nelson
Brig Gen Gil Hawk
Mr. Robert Beaton
Mr. Mike Livingston
Prof. Robin Murphy
Mr. Tom Wade

Migration
Mr. Thomas Saunders
Ms. Teresa Lunt
Dr. Scott Renner
Prof. V.S. Subrahmanian

Commercial
Prof Gio Wiederhold
Mr. Tom Clark
Mr. Gary Edwards
Mr. Scott Fouse

Viability
Dr. Duane Adams
Mr. Michael Brenton
Prof. Eugene Spafford

Support
Maj John Pernot
Lt Col Paul Schubert
Maj George Hart
Maj Mark "Ringo" Larsen

The study consisted of twenty-three people, drawn from government, academia and industry. Dr. James Hendler of the University of Maryland chaired it; Dr. Bhavani Thuraisingham of the MITRE Corporation and Mr. John Gilligan, the Chief Information Officer (CIO) of the Air Force assisted him. Included in the study were two general officers (one serving, one retired), six Ph.D.s and four serving officers who provided logistical and technical support.

Slide 2: Database Migration - The Goals of the Service

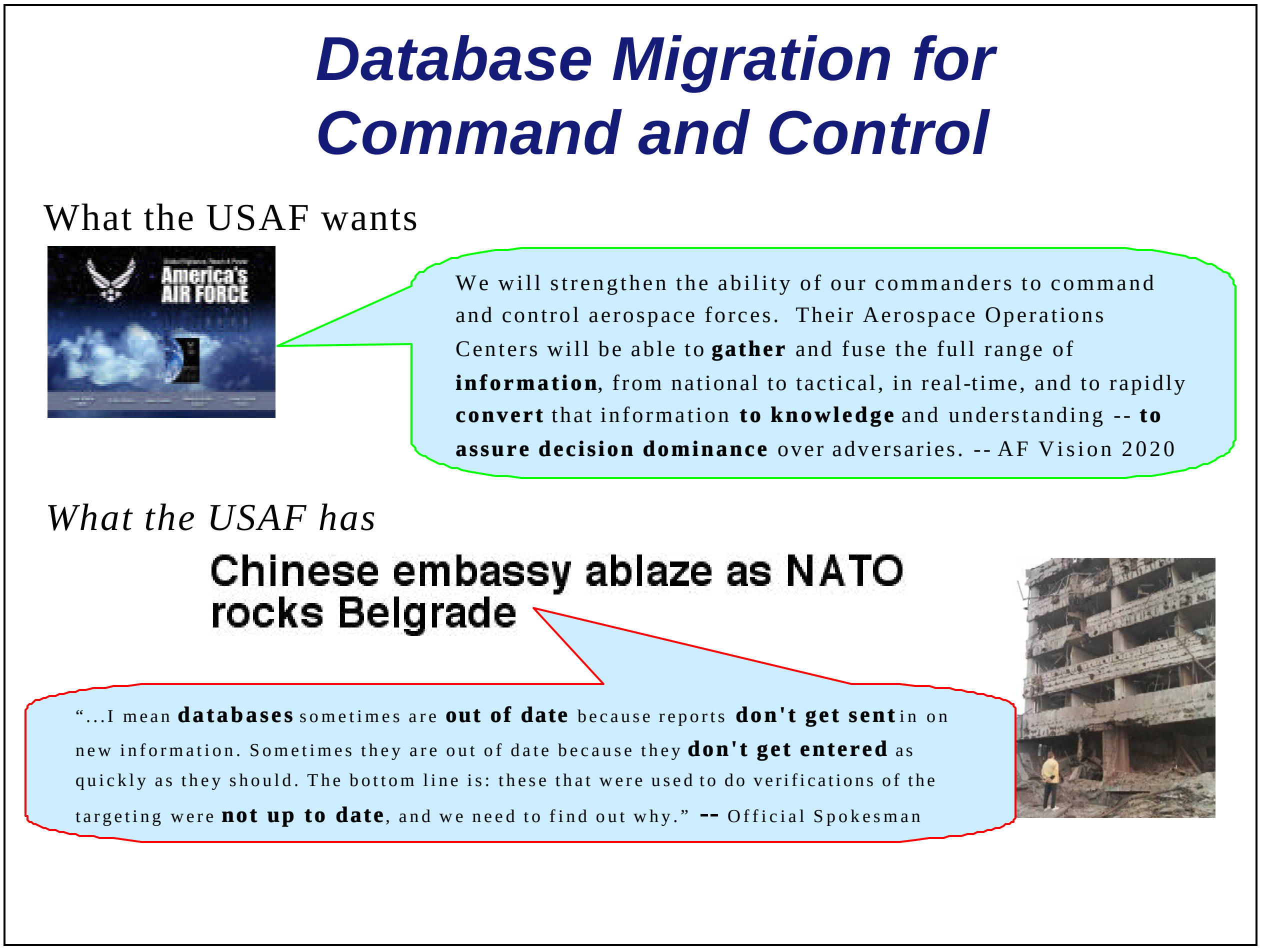

The United States Air Force (USAF) is striving to develop an integrated and seamless information management system for the command and control of Air Force (AF) operations. The goal, as described in documents such as *Air Force Vision 2020*, is to be able to have complete information dominance over our adversaries by the fusion of a wide range of data from multiple sources and systems.

Unfortunately, the reality of current AF systems is far from this ideal goal. Although our ability to execute operations is extremely good and our targeting is precise, data errors can cause major problems. Current databases are not integrated, are updated slowly, and are susceptible to inconsistencies that arise from manual input of data, a process unfortunately necessitated by the incompatibilities of applications, data stores and interfaces. In addition, troops are not trained to realize the importance of information sharing, and the updating of many different systems during the prosecution of operations is often of necessity lower in priority than the immediate and local decision making required by the command and control (C2) needs of the moment.

Slide 3: The Goal: Decision-Quality Data

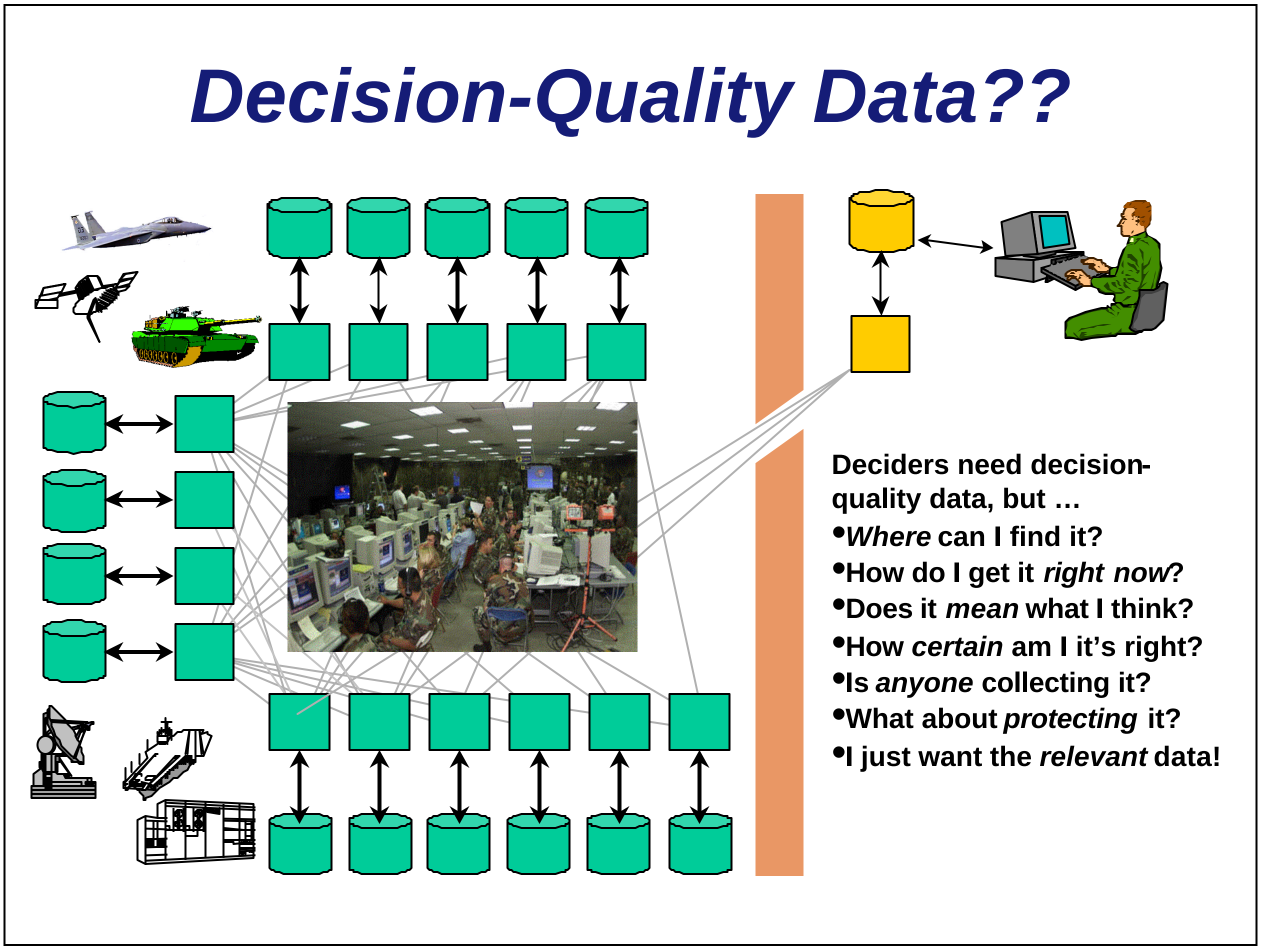

The C2 data resources in the AF are not currently able to satisfy the goal of supplying our warfighters with the needed decision-quality data. At present, these data resources exist as many separate "islands" of data. Each island is maintained by a separate organization for its own purposes and each island is largely unusable by others. These separate databases typically belong to individual systems, which communicate through expensive and inflexible pairwise interfaces. The effect, from the perspective of the decision-maker, is shown in Figure 4-3 (see page 79): Most data is hidden from the decider behind an impenetrable wall. He can obtain whatever data his system provides, plus a little data through the "peephole" built by his system developers, but if they have not perfectly anticipated his needs, then he will have no good way to obtain the missing data. The *best* he can hope for is that he might find some of that data by using another system, provided by another developer – leaving him with the difficult task of somehow integrating these separate facts *by hand* into a coherent whole. The wall between deciders and the decision-quality data has many facets, for example:

- *Discovering needed data:* Often deciders do not even know that the data they need exists. They may know it exists somewhere, but have no way of knowing where to retrieve it from, or whom they should query for needed information.
- *Accessing data:* Knowing where to find data is not always enough. The decider's system may be unable to retrieve data from the data resource for any one of several technical interoperability reasons.
- *Retrieving just the right data:* Sometimes, access to too much data is just as bad as no access at all. Deciders must have a way to describe the precise data of interest and a means to act upon that data.

In many cases, users get an overwhelming amount of data with an insufficient amount of time to process it into a usable form.

- *Fusing/correlating data:* Data retrieved from multiple sources typically needs to be combined into a single coherent whole before it can properly inform the decider. For example, in the targeting world there is often no way to decide whether data from separate databases describes the same real-world target, or to decide what to do when multiple descriptions of a single target conflict.
- *Obtaining accurate data:* When data does not correctly describe the real world, it isn't decision-quality data.
- *Having needed data collected:* There's no guarantee that the data needed by deciders will be collected and maintained by anyone. And there's often no way to effectively communicate the unmet needs to the providers best able to fulfill them. This problem is worst when there are system and/or organizational boundaries between data consumers and providers.
- *Knowing the data's provenance:* The data may be correct, but if the decider doesn't believe that it's correct, then it won't be used. Deciders can't judge the reliability of the data they receive unless they can know its source.
- *Knowing the data's timeliness:* Sometimes data becomes obsolete. Deciders can't be properly informed by obsolete data unless they know when it was entered, and when it was last refreshed or updated.
- *Having authorized access:* With most data, there are some people who should have it, and others who should not. Imperfect access control mechanisms and policies often keep data away from some legitimate users. However, the policies put in place to deal with this issue may also cause problems of their own- overly strict access policies are enforced to avoid leaks, for example, thereby keeping data away from some users who may need it. The more flexibly a database system can define and control access, the more "correct" the access policies can be.
- *Coping with multiple opinions on the same facts:* Many data resources only record a single "correct" version of the facts, even when multiple opinions exist and are important to deciders (this is not the same as data correlation, above - there we want to resolve multiple sources into a single opinion; here we want to put multiple opinions into a single source). An example of this may be a target for which multiple locations exist encoded in different databases. Unless a mission planner knows that there are multiple values, he is likely to act on the single value he retrieves, not knowing that elsewhere in the system there may be conflicting information that should be checked.
- *Understanding the data:* Semantics, or the "meaning" of data, is an element in almost all of the above problems. The meaning of data is something that comes from the people who work with it (or who write the software that processes it). When these different people have different ideas about the fundamental nature of the same data, it is very difficult to use that data to make good decisions.

Slide 4: The C2 Data Problem

The C2 Data Problem

- AFSOC (Hurlburt AFB)
  - 10-15 DBs for night-time helicopter missions
  - 5-10 people dedicated to manually translating between DBs during mission execution
  - Training missions cancelled or range time lost at least 1X a week

- AF/IL (SSG, Gunter AFB)
  - 120 systems, 2000 interfaces (30-40% of all code)
  - 39 servers for LOGMOD, firewall management nightmare
  - Data standardization (3 ILM systems) cost $40M, 4 years

- 7th AF (Osan, Korea)
  - TBMCS support to Integrated Tasking Order (ITO) Preparation
  - Facility target datasets failed to load (over 8,800 discrepancies)
  - ITO delivered later than required
  - Development of local work-arounds - Separate "off-line" database for aimpoints

To determine the operational impacts and importance of these database interoperability problems, the panel visited a large number of sites and surveyed many different databases and the systems where they are deployed. We explored the impacts on operations at a wide variety of scales, ranging from relatively small missions such as Air Force Special Operations Forces (AFSOF) night training up to the large-scale operations of the Numbered Air Forces. In all these cases, we saw or heard about problems similar to those shown above. The impacts of these types of problems are quite high.

Slide 5: The Problem

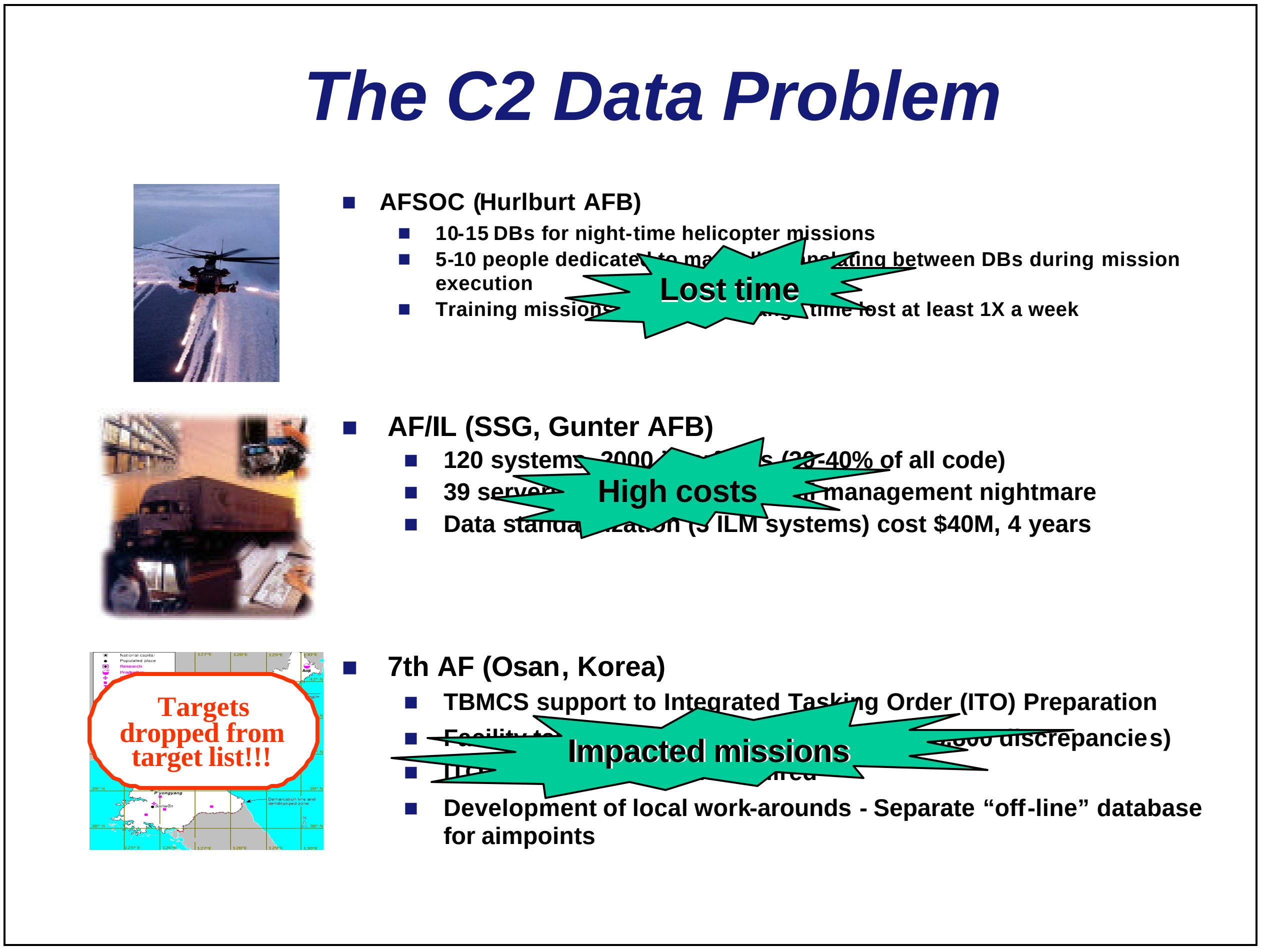

Any understanding of the problem will proceed most usefully from a study of some examples. AFSOF has a number of airmen and women who must spend their time supporting databases for missions to be successful. Often, however, especially when the support is for training, these valuable people have higher-priority tasks to pursue. The inability to automate database intercommunication forces commanders to make a difficult choice: either these people must be pulled off important tasks, or training must be cancelled. The impact is high - as the size and complexity of operations grow, the number of people wasting their time increases correspondingly, and the scale of the problem grows.

A particularly egregious case of this can be seen in large support units, where the number of people needed comes to exceed the unit's ability to provide them. The number of custom interfaces grows, and the support for these interfaces (estimated at 30-40% of the code in some cases) consumes vast portions of the Operations and Maintenance (O&M) budgets of the units -- an unacceptably high price for these organizations. Furthermore, trying to develop "post hoc" solutions to integrate systems developed separately using a plethora of different architectures, commercial off-the-shelf (COTS) systems and versions of code costs a significant amount of money and imposes burdens on these units that prevent them from modernizing systems and upgrading capabilities.

Similar problems also show up in the major command (MAJCOM) C2 systems. The interoperability problems lead to duplication of effort, waste of money, and operational impacts as inconsistencies can grow in the various versions of the systems. The operational impacts of these system problems can lead to major "errors" by the Air Force. The results can include the prosecution of improper targets, a problem that puts pilots at risk and leads to other command and control problems with far-reaching impacts.

Slide 6: Getting To The Present

How Did We Get Here?

- **Operator functions were the focus of C2 systems development**
  - **Mission success required collection of significant amounts of data (e.g. targeting, planning, scheduling, …)**
- **Data to support the functions has been part of the systems**
  - **Early database technology created rigid and hard-to-extend databases**
  - **Users had to build interfaces and local database copies to maintain correct local data values or extend data models**
    - **Interfaces proliferated, raising maintenance costs and support nightmares**
    - **Multiple data entry is needed increasing error rates**
    - **Data entry is low priority under operational pressure**

Our legacy problem will require significant effort to overcome

The Air Force's problems have been a long time developing and will take significant effort to overcome. Later in this report (see page 79), we describe the technical problems in greater detail. This summary slide shows that the way we develop and build our systems causes far-reaching problems that will not be easily overcome. Although the technology available for database interoperation has improved dramatically in recent years, the policies and procedures we have for developing systems have not changed, and the problems we've discussed so far continue to increase and propagate.

Slide 7: Where We Need To Go

Where We Need To Go

Command and Control
(From "Transforming the Military")

Decision support: The ability to develop accurate estimates of future enemy and friendly capabilities

- **Real-time ISR support to combatants at all levels**
- **Integrated all-source information in near real-time at each operating headquarters**
- **Constant, accurate common relevant operating pictures to include the lowest operating entity**
- **Cross-service platform and echelon integration of relevant information**
- **Integrated management of intelligence collection with the planning and operations cycle**

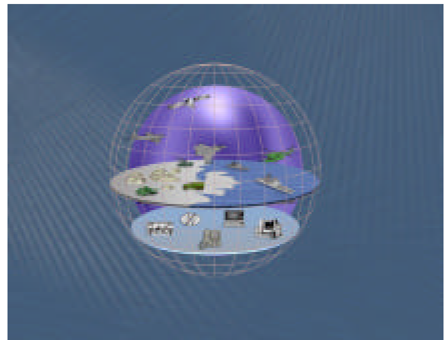

Data fusion is enabled by agent-based computing and an information infrastructure that automates significant data acquisition efforts currently performed by hand.

Future C2 CONOPS critically dependent on high-quality fused data

The USAF is particularly concerned because the increasing complexity of Air Force operations requires increasingly complex information integration. The concept of operations (CONOPS) emerging from other Scientific Advisory Board (SAB) studies (such as the soon-to-be forthcoming SAB study on *Predictive Battlespace Awareness to Improve Military Effectiveness*) show the need for an information infrastructure that can automate fusion and provide far more capability with respect to decision-support needs. In short, we face a need for better and better integration, but our systems are not increasing in capability so as to support these needs, a problem that we must exert significant efforts to overcome.

Slide 8: Can't Get There From Here

Can't Get There From Here…

- **Today's C2 missions demand greater information sharing than our current data systems can deliver**
  - **Databases are not shared and thus fusion into information and knowledge is not happening (human or automated)**
- **We acquire systems which don't facilitate data interchange**
  - **Integrators are allowed to control the database as a private component of their system**
- **Data interoperation is added as an afterthought**
  - **Too costly and too time-intensive to implement the required data sharing**

…without a significant change in how the AF manages C2 data

Summarizing the points raised to date, C2 missions demand data interoperability at a level not supported by current systems. The way we build and acquire our data systems is not facilitating the increasingly complex needs of the modern Air Force, and we cannot simply continue to add data interoperability as an afterthought. The costs, time, and personnel requirements are beyond the capacity of the Air Force. If we don't make significant changes in the way we manage our C2 data, we will never be able to achieve the operational effects and improvements the Air Force will need in the future. The remainder of this report describes how the USAF can overcome this problem, and describes processes and procedures that can help us to reform our methodologies and reach our operational goals.

Slide 9: Terms of Reference

The Terms of Reference for our study are shown in this slide. The study was asked to recommend methods for the migration of current databases to a modern information infrastructure. In addition, it is clear that the current legacy systems cannot simply be turned off while we wait. We need to be able to continue to keep a viable data infrastructure at all times as we migrate to more modern approaches.

The term “migration” has a number of meanings and connotations. Appendix A is a more detailed discussion of the many issues that complicate database migration efforts.

Slide 10: Study Methodology

Study Methodology

- **Surveyed USAF C2-related DBs and their acquisition to understand the current situation**
- **Surveyed migration efforts (Commercial and DoD) and support technologies to identify key success factors**
- **Explored how the USAF can apply these lessons in migrating our C2 DBs within the contexts of our specialized culture and missions**

We organized our study to explore several issues simultaneously. We needed to understand both the current situation and the forces that keep us from overcoming the current problems. We surveyed trends in the commercial world, examining both successful and failing migration efforts to see if we could identify key lessons learned there, and we have thought about how these lesson could be applied within the specialized context of Air Force Command and Control operations. Finally, we will develop a set of recommendations by which these processes and procedures can be put in place.

Slide 11: Outline

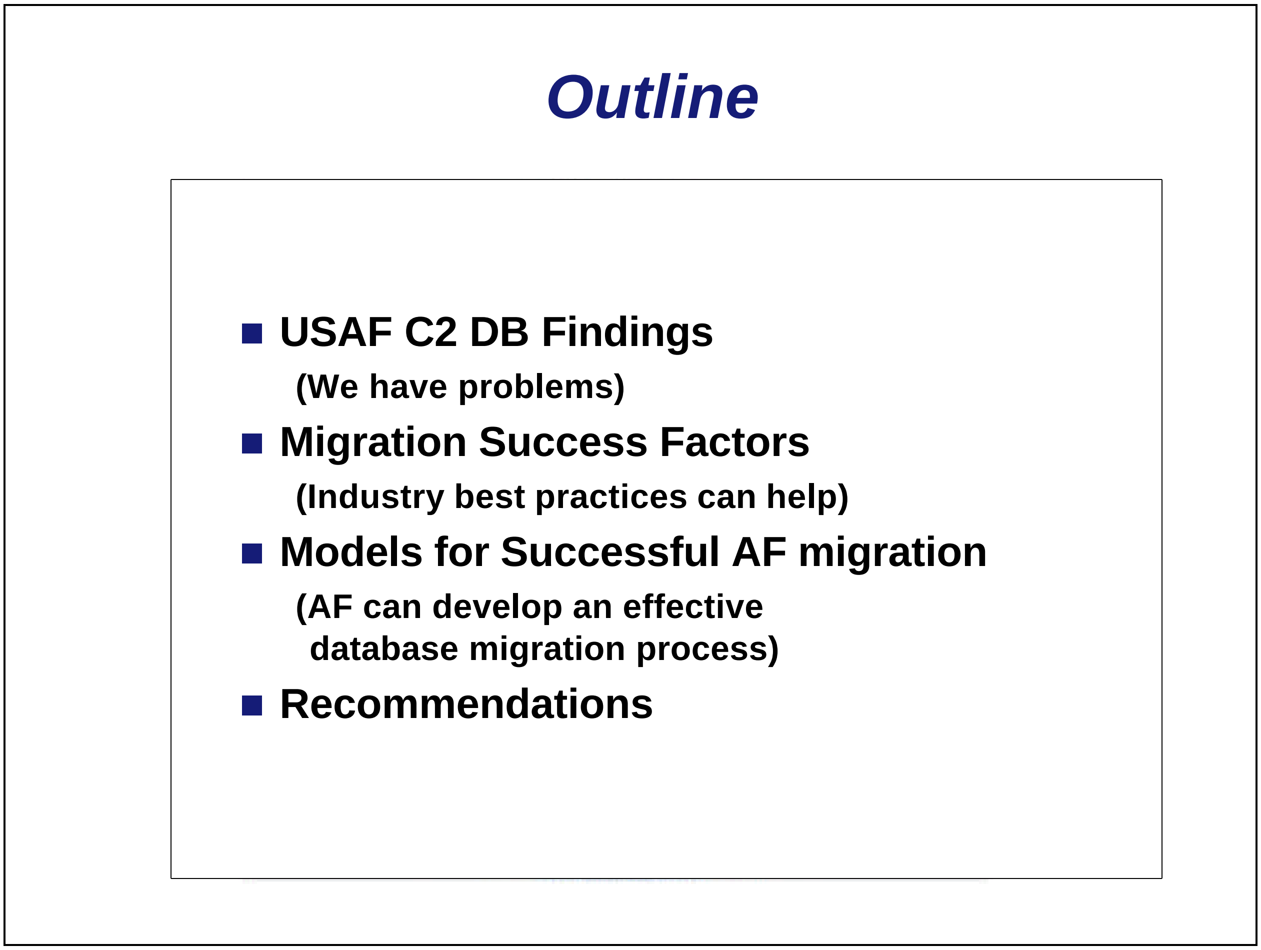

The remainder of this report will have four major sections:

- First, we present what we have found in our analysis and exploration of current Air Force C2 Database systems.
- Second, we summarize what we have found that works in the context of successful commercial database migrations.
- Third, we briefly describe some Air Force-specific steps that might help us to get our data migration started.
- Finally, we outline specific recommendations that must be taken for the Air Force to solve its data migration problems.

Slide 12: Findings

Current C2 DB Findings

In this section, we summarize current AF database systems and the acquisition processes that cause our data interoperability problems. This section presents a set of slides that summarize the current situation and the problems it causes. Following the slides we present a technical section, which describes current acquisition practices and problems in great detail.

Slide 13: Databases and Systems Surveyed

Systems/DBs Surveyed

- TBMCS DBs
  - AODB
  - EMRDB
  - TAPDB
  - Track DB
  - TCT DB
  - IDPDB
  - TWM
- TBMCS-UL
- AMC DBs
  - GDSS
  - C2IPS
  - GTN 21
  - CAMPS (CMARPS-ADANS)
- EUCOM – VAMP
- STRATCOM
  - Enterprise Database (EDB)
- PACOM – CHP
- Intel
  - MIDB
  - IPL
  - 5D
  - MEPED
- Logistics
  - DCAPES
  - JOPES
- SPACECOM
  - SPADOC
  - CCPDS-R
  - Granite Sentry
  - ASW
- GCCS
- GCSS
- USAFE – RAMP, COIC-TDB
- MSG Data Depot
- GTN
  - GATES
  - JALIS
  - CAMPS
  - RCAPS
  - GOPAX
  - CFM
  - DTTS
  - DTRACS
  - WPS
  - ICE
  - IBS
  - CMOS
  - TC ACCIS
  - DAAS
  - MTMS
  - RF Tag
  - CEDI

Survey included several "generations" of AF DB systems

Our study panel examined a large number of the databases and systems used in Air Force Command and Control. We explored systems that directly supported operations (such as the Theatre Battle Management Core System [TBMCS]), databases used in the support of C2 (including logistics and combat support), and systems used in some of the functions (such as intelligence) that are critical to C2 mission success. We also completed a detailed survey of current acquisition approaches and their effect on the systems in development.

We present our findings in three separate ways. First, we summarize our findings as a set of brief bullets and lessons learned. Second, we summarize the current acquisition system to show how the systems are developed with respect to requirement setting and funding. Finally, in the "Current C2 Database (DB) report" section, we provide a comprehensive overview backing up these summary findings.

Slide 14: Findings

AF DB Findings

- DB problems plague AF systems, old and new
  - Manual replication of data
  - High maintenance costs
  - Duplicate, stale copies of DBs
  - Structure of DB inhibits change, evolution
  - Information access/security methods impractical
  - Proliferation of “local” DBs to work around shortcomings of systems of record

New database use not showing significant improvement over old

The most important finding of our study is that we do not see a lot of differences between systems acquired in recent years and those of an older vintage. Where we look at new database systems we see all the problems we have identified (and discussed previously) in older systems. The systems include significant needs for manual data entry in multiple places, high maintenance costs, and all the other problems shown in this slide. We find that the operational effects described earlier are being seen throughout the USAF C2 enterprise.

Slide 15: Findings

AF DB Findings

- **AF MAJCOMS developing operational architectures for C2 missions (AFSPACECOM, AMC, ACC)**
  - **Information Exchange Requirements just beginning to emerge**

Integration across these efforts (CAF C2) still lacking

- **Some new database efforts moving in the right direction**
  - **AFSPACECOM, USTRANSCOM (AMC), AF/IL - applying successful commercial business processes**
  - **AF Portal - exploring new management techniques for "corporate" AF data**
    - **Lessons learned will be of use to development of needed C2 portals**

No coordinated effort to share lessons learned or migrate best practices

We did see encouraging signs in some of the ways that current systems were using COTS databases and database components in their systems development. The MAJCOMs have been tasked with the development of operational architectures for C2, and a number of them have begun to define their C2 CONOPS. To date, however, very little integration of these efforts has been seen, and the Information Exchange needs across the enterprise have not been carefully developed, analyzed, or operationally tested.

In addition, we do see commercial best practices starting to influence some Air Force database systems. Business processes used in industry are starting to influence the Air Force system construction, and the maintenance costs and support needs are consequently improving. In addition, we are seeing the development of portals, most particularly the Air Force portal that is being developed to improve data management in the corporate Air Force.

However, the coordination between these efforts is still informal and the "best practices" are not often shared. Each organization makes its own decisions on matters such as COTS tool selection, integration approaches, and data modeling. The Air Force is thus unable to take advantage of economies of scale or to encourage integration across the entirety of the Air Force C2 enterprise. It cannot promote coordination between the various MAJCOMs or the various functions within the Air Force that support C2 operations.

Slide 16: Findings

AF DB Findings

- **Data sharing is a hallmark of successful unit-level developments**
  - **e.g. Web-enabled USAFE target folders used in Kosovo operations**
  - **e.g. PICCS (TBMCS-UL) improvement in dataflow among units**
- **Acquisition model for unit level efforts still evolving**
  - **Nascent “UDA Man” program at ESC developing “incubator” approach**
  - **CAOC-X considering fielding of unit level efforts**

Mechanisms to support unit development efforts needed

Another encouraging sign is the sight of the kind of innovation in Air Force units that allows them to prototype their own systems and to improve coordination between the units. A particularly noteworthy example of this is the Pacific Air Forces Interim Command and Control System (PICCS) effort developed at PACOM, now going into acquisition as TBMCS-Unit Level. Another example is a target folder application developed at USAFE during the Kosovo operation, which provided web-enabled target folders. However, these unit level development efforts are not widely recognized as beneficial within the acquisition community and the operational military. Some organizations are starting to provide support for these efforts, but a more coordinated approach is necessary. In addition, successful unit development efforts such as PICCS become saddled with additional duties and responsibilities once systems they develop become of use to the Air Force, providing a strong *disincentive* for units to publicize and share their useful software solutions.

Slide 17: Findings

AF DB Findings

- **Emphasis is still on data collection, not data sharing and fusion**
  - **Central leadership not focused on AF enterprise C2 data**
  - **Consistent management structure not in place**
    - **MAJCOMs differ markedly in data management approach/capabilities**
    - **Disconnect between Unit (bottom-up) and CIO (top-down) data integration efforts**

Data is not currently managed as a C2 enterprise asset

One of our key findings is that the Air Force does not have working mechanisms to foster the sharing of data across the C2 enterprise. Data and databases are generally seen as falling simply within the purview of the "stovepiped" acquisition processes of their owning organizations, and the MAJCOMs and Air Force functional organizations get to define system level requirements without any examination of the overall effect across the whole C2 enterprise. Central leadership is beginning to recognize the importance of this, but many in the command structure still do not recognize it as a critical need. Furthermore, the shortcomings of the major systems deployment process puts a pressure on units to develop their own solutions, but these solutions often cannot use the same standards that top-down systems are trying to mandate. In addition, these units must often pay for their own licenses and maintenance, essentially paying "retail" for what the larger Air Force enterprise could get "wholesale." For example, many units are forced to buy separate licenses for Oracle or other database products that could be provided more cheaply if the Air Force had licenses that spanned the enterprise. In addition, the units have as their primary focus the development of software that meets their needs, rather than the production of a data stream that would be accessible to higher command or other units across the enterprise. This culture of local fixes to acquired systems again reduces incentives for the sharing of software and data across the Air Force C2 enterprise, and creates major problems for the Air Force in the prosecution of time critical targets; it can also contribute to costly errors that arise as products of inconsistent or wrong data.

Slide 18: Findings

AF DB Findings

- **DB-admin/support staffing problems are critical across the AF**
  - **Not just career path/retention issues**

Current system development exacerbates, rather than alleviates, DB admin staffing issues

- **Significant problems with data interoperability, term agreements, data dictionaries, etc.**
  - **XML being viewed as silver bullet (It isn't)**

AF under-invested in S&T for interoperability

Although it is clear that there are career path problems and retention issues that effect all information technology workers in the USAF, there are particular problems for database administrators and systems administrators that are caused by the poor way we build and field database systems. As newer systems enter the command environment, they are usually built to different standards and often run using incompatible software products or server platforms. Interfaces between these systems and current systems must often be built to overcome these incompatibilities, and they add a training need: administrators who know how to run the current configuration must be taught to maintain yet another data stovepipe. Building to a common set of products and standards could help to alleviate this problem, but instead Air Force culture forces an increasingly complex task on the already overtaxed administrators. The increasing burden on the support staffs (and on the O&M budgets that pay them) is running at direct odds with the Air Force's needs to reduce this sort of staffing and to lessen the burden on the existing staffs.

Another problem we observed is a belief that somehow the commercial world will wave a magic wand, utter the incantation "XML" and magically improve everything. Unfortunately, while modern technologies such as the "eXtensible Markup Language" provide important tools for helping to fight interoperability problems, they offer only part of the solution. These technologies provide some of the "plumbing" needed to help fix our data problems, but they leave major issues unattended. They do not help us to bring our current legacy systems together, they don't provide help in the development of the data dictionaries and term libraries (ontologies) needed to make our systems compatible, and they fail in and of themselves to provide standard services across the Air Force. Processes and systems to use these approaches are being developed within the Department of Defense (DoD) research community, particularly at the Defense Advanced Research Projects

Agency (DARPA), but the Air Force is significantly under-invested in the business of transitioning these technologies and adapting them to Air Force needs. We discuss these issues further later in this report, but one of our key recommendations is worth stating forcibly at this point:

- The Air Force should significantly increase the funding of the Joint Battlespace Infosphere (JBI) project at the Air Force Research Laboratory's Information Directorate (AFRL/IF) -- this is the Air Force group doing the most important interoperability S&T in the enterprise, and at the current time it is inadequately supported and undervalued.

Slide 19: Findings

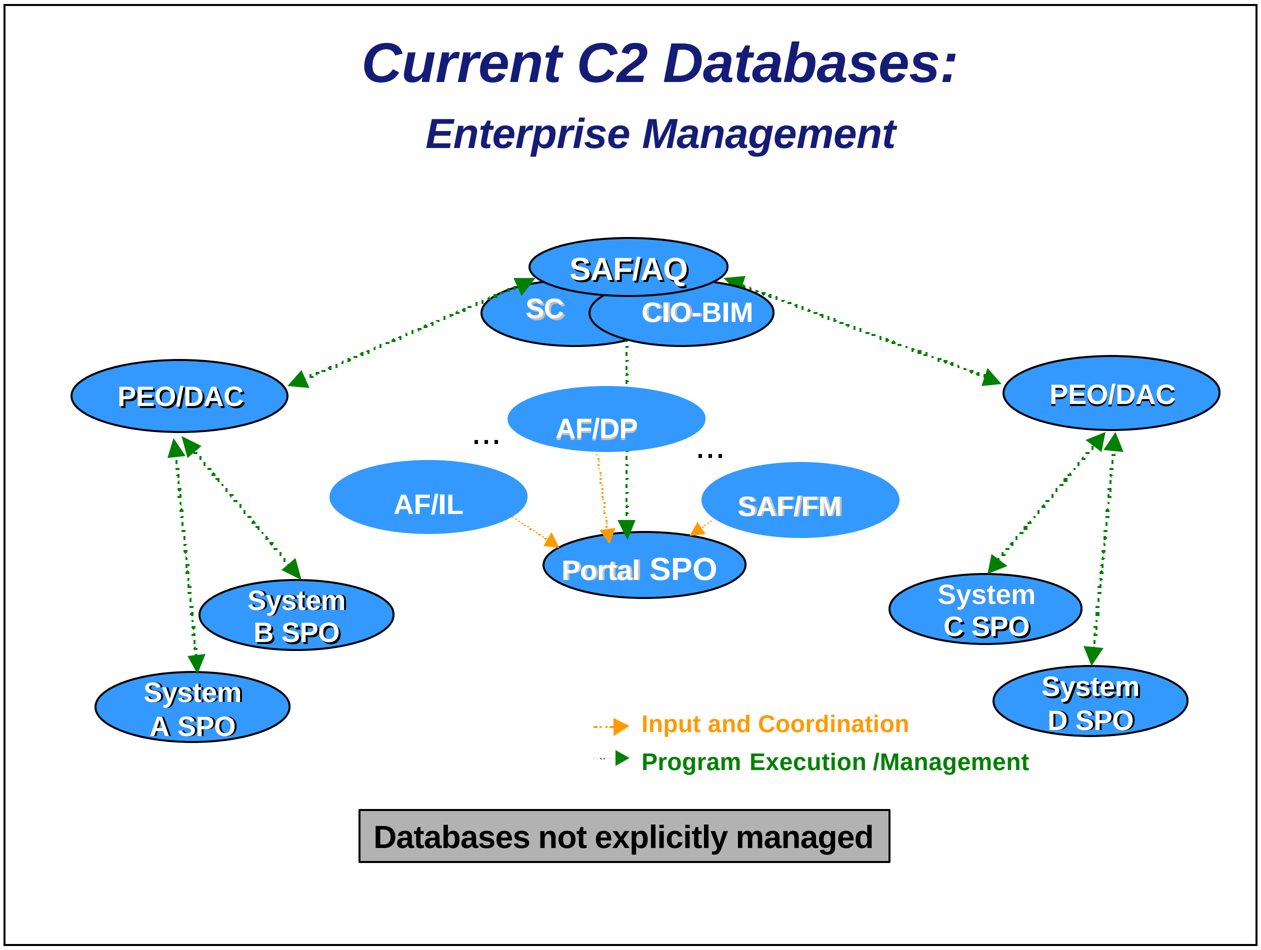

In this slide and the next we summarize the situation with respect to acquisition and development of C2 database systems in the current Air Force construct. We focus on management across the enterprise and on the requirement flow. Later in this report (see page 103) we will contrast the current approach with the approach we advocate; our approach will enable the Air Force to take immediate and effective steps to start overcoming this critical C2 deficiency.

There is essentially no C2 database integration construct being used in the USAF. Individual programs characterize the manner in which their information is to be kept and handled. Where interaction with other systems is necessary, each system is left to incorporate procedures to establish its internal interface rules and to conform to whatever pair-wise information exchange processes have been established with interfacing components. Databases are not explicitly managed.

The only disciplined requirements processes, execution processes, and coordination processes are those that relate to the traditional style of defining information flow across system boundaries embodied in Interface Control Documents (ICDs). Information initiatives such as the Air Force Portal are attempting to influence better behavior, but standard forms of expressing information have not yet become accepted, and the content of information to be displayed via portals has not become the subject of agreement. In an effort to move forward with portal technologies, the Air Force CIO should identify information stewardship by mission area. Mission area councils will identify sub domains; each subdomain will nominate the data to be used (i.e., exposed) beyond their domain. Although steps in the right direction are being taken, progress in this area remains slow.

Slide 20: C2 Database Requirements

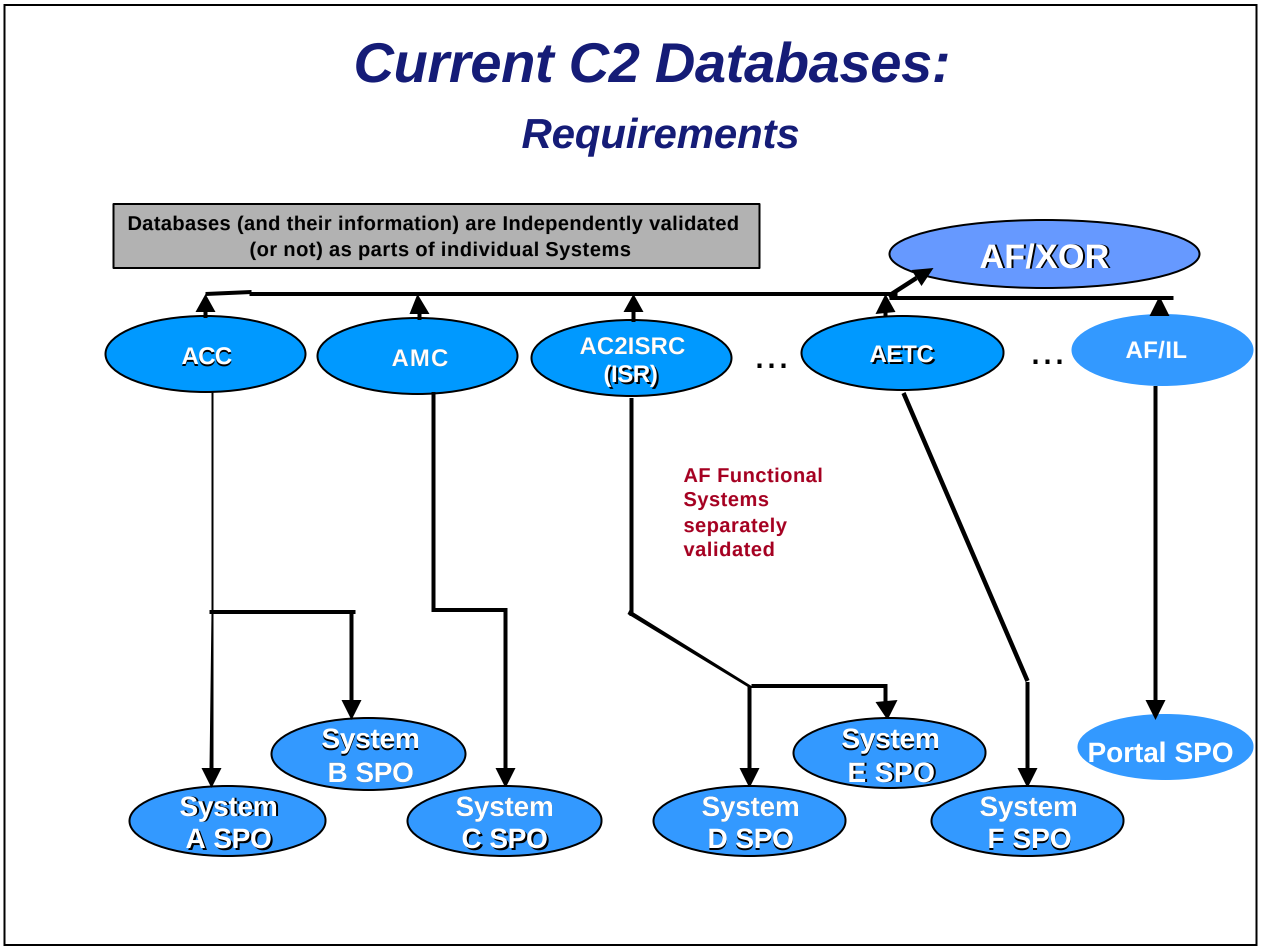

Much of this problem can be attributed to the current requirements-driven approach to acquisition used in information systems acquisition by the Air Force. The requirements are generated locally and are not coordinated across the Air Force C2 enterprise. In the current information management construct, individual platforms develop both database and information exchange requirements for approval and validation by the MAJCOMs and the Deputy Chief of Staff for Air and Space Operations' Operational Requirements Division (AF/XOR). The individual platform programs allocate these requirements down to the individual platform acquisition System Program Offices (SPOs). The SPOs negotiate amongst themselves for the manner in which information might be shared, but no systemic solution has been visualized or adopted. The C2 Database problem is a result: there is no requirements process for managing C2 databases.

Slide 21: Findings Summary

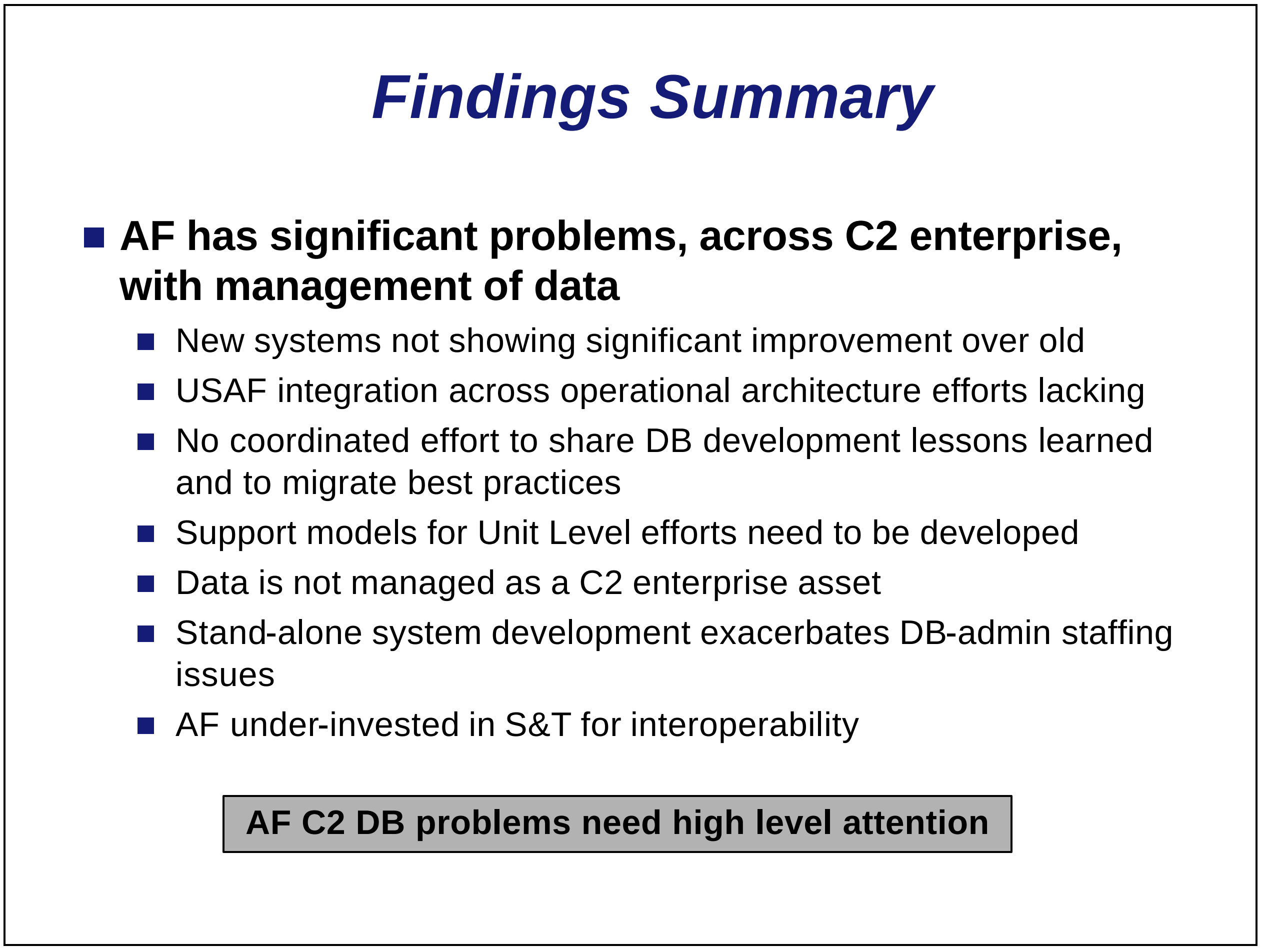

Summarizing our key findings - although there are some encouraging developments, particularly at the unit level - the overall C2 data problems of the AF are *NOT* being caused primarily by technical deficiencies, but by acquisition and cultural problems that cannot be solved by technological magic. The Air Force needs to assign high-level attention to this critical problem, making enterprise-wide integration and sharing of data a high priority item. Science and technology (S&T) work in this area is under-funded, and the system development is not only failing to fix the problems, it is making many of them worse. While some in the current Air Force leadership do appreciate the importance of the data integration problem, there is not yet a general awareness of the importance of a solution for this problem or of the need for a plan to address the current shortcomings. These shortcomings will be discussed in detail in Chapter 2.

(This Page Intentionally Left Blank)

Chapter 2

Current Air Force C2 Systems - Databases and Migration Plans

Slide 22: Current C2 DB Report

Current C2 DB report

2.1 Summary

2.1.1 Background--The Current AF C2 Acquisition Process

The Air Force has for some time acquired and sustained its Command and Control systems (and established and evolved their concomitant databases) using much the same process it uses to acquire major hardware systems such as the F-22. The process is one that presumes that the requirement for a new system is a product either of the obsolescence of existing systems or a major technological breakthrough. The current technology is then evaluated and used to develop a performance specification, which is included in a Request for Proposals after an appropriate budget line has been established. The resulting contract usually runs over a period of years and has a value of a few hundred million dollars. Development and operational testing is accomplished when appropriate. There are usually major milestones specified during the development and testing period, when the Air Force or DoD Acquisition Executive is required to certify that proper progress is being made, or (in its absence) to restructure or terminate the program. There is no production phase involving significant funds as there is with major hardware programs. After successful operational testing, the developed software is replicated for the field sites and new computer and communications hardware is procured to support the new software. Initial Operational Capabilities (IOC) and Full Operational Capabilities (FOCs) are normally defined. The system then enters a sustainment phase, which is marked by lower levels of annual funding than the acquisition phase, and improvements in operation and support are realized through O&M funding. This process is described in the DoD 5000 series directives, is generally used by all the services, and is sometimes known as "the big bang" process, because it generally starts from a clean sheet of paper and involves the creation of a new product in a single discrete contract, rather than the evolution of a new system from an old one; a single contractor usually has the responsibility for such work, a responsibility that has grown with the downsizing of the acquisition corps.

2.1.2 Data in the C2 Systems

As each new system is brought into operational use, operators are retrained to use the new hardware and software. The new contractor usually does this retraining. The data, where it is similar to the data in the existing system, may be retained and represented in similar form. But regardless of the circumstances, the data in the system is viewed as a part of that system, not as an important quantity in and of itself. This happens even though the same data may be in use elsewhere. There is normally no explicit constraint on the developing contractor to retain operator processes or to establish multi-system databases, when the same or similar data is to be used by a number of systems. This latter concern is beyond the authority of any individual system development contractor.

2.1.3 Consequences of the Current Acquisition Process

Since each C2 system has been developed as a separate entity, and interfaces to other systems are controlled through interface control documents and various messaging systems, we have evolved a process whereby operators maintain similar data sets in various systems across the Air Force. Each of those systems requires a separate set of update and error-checking processes, with the concomitant duplication of resources (people, time, and money). The consequences of these procedures are exceptionally wasteful, both in terms of resources and in terms of effective operational capability. A much more efficient approach, in terms both of the use of resources and of operational effectiveness, would be to focus our efforts on the timely and error-free maintenance of the data. Various Air Force C2 systems could then access that data as needed for their particular functional requirements. This, of course, is the basis of the JBI concept. The current approach to funding and executing our acquisition programs precludes that solution.

2.1.4 C2 System Acquisitions Demonstrate Inadequate Planning for Sharing of Data

We examined many Air Force, Army, Navy, and joint C2 systems during the course of this study. By way of exemplifying the data issues extant in those systems, we discuss in the following paragraphs three system acquisitions and collections of system acquisitions: the Air Force Space Command (AFSPACECOM) Integrated Space C2 (ISC2) effort, TBMCS, and Air Mobility Command's (AMC) component of U.S. Transportation Command's (TRANSCOM) Defense Transportation Enterprise (DTE).

2.2 Current Air Force C2 Systems Reviewed

2.2.1 The Integrated Space Command and Control System (ISC2) Contract

2.2.1.1 ISC2 Introduction

The Air Force let a contract to Lockheed Martin in the fall of 2000 to evolve the Command in Chief of North American Aerospace Defense Command/Commander in Chief of US Space Command's (CINCNORAD/ USCINCSPACE) ISC2 system. The ISC2 effort is defined in an evolving Target Operational Architecture (OA)[1] developed by the using organization(s). An evolving Target System Architecture (TSA)[2] has been developed by the Contractor in response to that OA. The TSA addresses key elements of the NORAD/USSPACECOM Warfighter Support System (N/UWSS) Statement of Objectives (SOO), specifically focusing on the N/UWSS Users' Integrated Command and Control (C2) Vision, the Mobile Command and Control Centers (MCCC) Consolidated Modernization Roadmap, N/UWSS System Architecture Objectives, the N/UWSS Technical Architecture, the CINC C2 Node CONOPS, and the System Maturity Matrix (SMM). Since the ISC2 contract award, the DoD has changed the structure of the component commands, creating (for all intents and purposes) a separate (4-star) service from the Air Force Space Command, a major component of NORAD and US Space Command. The full impact of this reorganization is yet to be realized (interviews with ISC2 program staff in the Spring of 2001 produced the opinion that little would change with respect to program prognosis). With the creation of another 4-star command, there is a high likelihood that another "kingdom" will be established, with its own system acquisition and proprietary data.

2.2.1.2 ISC2 Background

The current collection of semi-automated C2 support systems has been used by NORAD (and its Air Force Components, Air Defense Command [ADC - now defunct] and Air Combat Command [ACC]), and more recently USSPACECOM (and Air Force Component AFSPACE), both of which are principally located in Colorado Springs, for over 40 years. The initial mission was continental air defense, which was augmented (and supplanted in importance) by missile attack warning when intercontinental ballistic missiles (ICBM) were introduced into the Soviet force structure. Since there has never been a missile defense system fielded in the US, we have depended on a national strategy grounded in the philosophy of "mutually assured destruction" (MAD), which was introduced in the 1970s. This philosophy has required "lots of nines" assurance of detection of impending missile detonations, and the ballistic missile attack warning aspect of NORAD/USSPACECOM's mission has been at center stage since the threat surfaced. Computational hardware and software has been and is now being used to maintain the "satellite catalog", and thousands of satellites (mostly debris) are routinely tracked. The continental air defense mission continues. AFSPACE also operates the US ballistic missile force, but the C2 aspects of that mission do not appear to be part of the ISC2 effort. The Air Force Electronic Systems Center (ESC) has been the acquisition organization for all the recent major C2 acquisitions. In the 1980s a number of acquisition programs which were improving the Space Surveillance and Tracking, North American airspace surveillance and Control, Communications, and Missile Warning capabilities of CINCUSSPACE/CINCNORAD were combined into the Cheyenne Mountain Upgrade (CMU) Program. After the painful recognition that integration of ongoing programs requires more money and time, not less, the $1.5B CMU program resulted in an operational set of federated systems in the mid-1990s. The evolved version of this and similar independent system developments now forms the de facto ISC2 baseline. Significant costs for sustainment and shortfalls in capability to address the needs of the post-Cold War environment led to a need for a new approach, and the ISC2 contract was let in 2000.

---

[1] See the C4ISR Architecture Framework Version 2
[2] Ibid.

2.2.1.3 ISC2 Recent Environment

2.2.1.3.1 Current ISC2 Output Users

The ultimate user of the Missile (and Air) Threat Warning Data is the National Command Authority (NCA). USCINCSPACE/CINCNORAD are required to assess the extent and severity of the attack. While there is no active defense against ballistic missile attack, NORAD does play a role in responding to the attack, with ACC being the AF defense component. It appears that a primary user of the gross space surveillance data is the National Aeronautics and Space Administration (NASA), as they are concerned about the safety of the Space Shuttle and Space Stations. US military and intelligence agencies are concerned about specific surveillance and communications satellites.

The attitude of the USSPACECOM/NORAD staff and command structure is that vital national decisions are based on the correctness of their information, and so there is an institutional tendency to hold critical data closely so that CINCNORAD or USCINCSPACE can be assured that he is undoubtedly conveying the correct picture to the NCA.

2.2.1.3.2 Scope of the ISC2 Effort

In the ISC2 effort, Lockheed Martin is involved in a new and long-term alliance with the AF, one without the traditional IOC and FOC or "big bang" development approach, but with a level of effort to sustain the current system and evolve to a new system architecture implementing a "to-be" Operational Architecture synthesized by the operators during a two year period before the ISC2 contract was let. There are annual program reviews by the "Milestone Decision Authority" (the "milestones" are not the traditional performance milestones) that approve the program of work for the coming year vice making cancellation or restructuring decisions on the program. This is truly a new paradigm for AF C2 acquisition, and the viability of the approach is yet to be demonstrated (the contract is roughly a year old). At present the level of funding is approximately $100M a year (FY 01), divided into 5 separate Program Elements (PEs); O&M consumes 2/3 of the funding. About 30% of the available funding is levied against sustainment of legacy systems. The $100M is programmed for sustainment and improvement of the baseline CMU (and other) legacy efforts. There is a twice-yearly update/improvement/correction of the legacy systems, and new system architecture is introduced as funding is available through a one-year spiral development process (a process that encompasses several minor spirals). This will be a difficult effort, because the budgeting/funding structure in the AF is not friendly to "wedges" of funding for future efforts when their content cannot be clearly defined—and dollars that are not identified against "legacy" systems (e.g., for sustainment) are consequently at risk. This vulnerability was already demonstrated during the FY 03 budgeting cycle, when projected savings were taken from the AFSPACE budget but only half of the capital to generate those savings was included (not an unusual feature of the budgeting process). The current concept of the ISC2 Enterprise is shown in Fig. 2-1 (see next page).

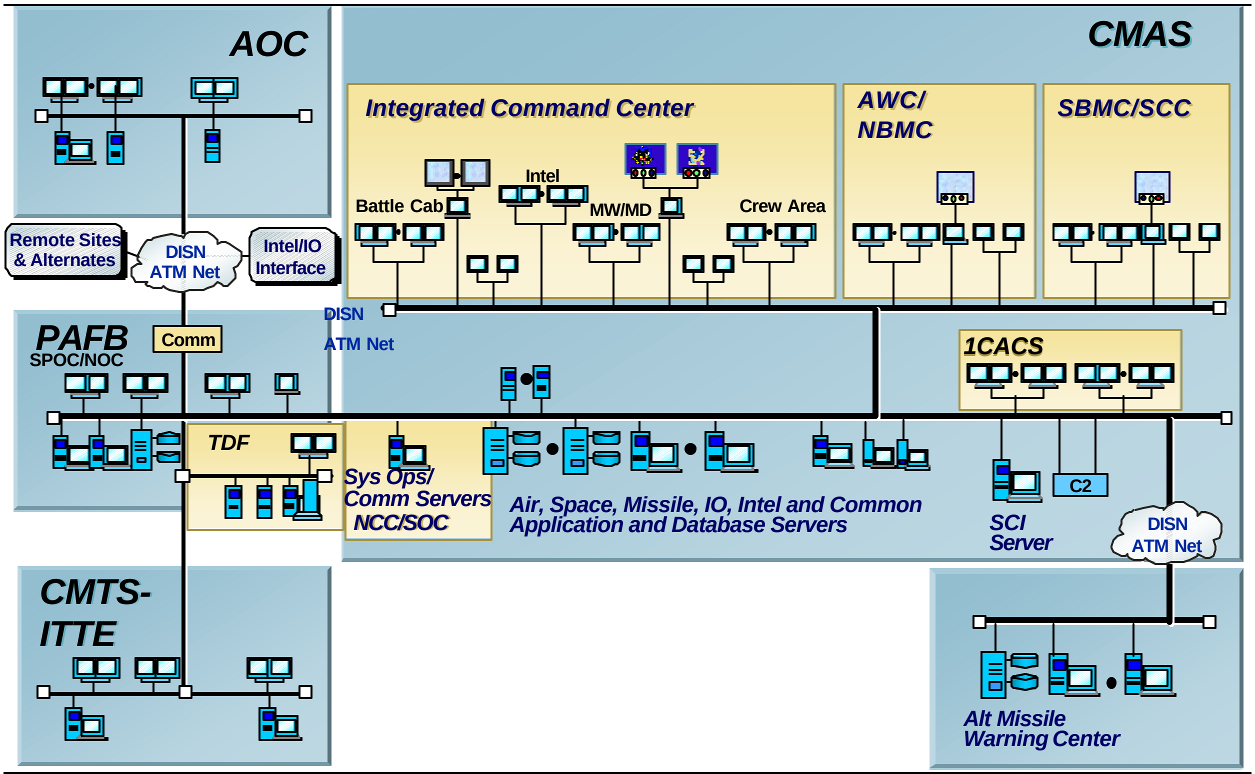

Figure 2-1. *Current Concept of the ISC2 Enterprise*

2.2.1.3.3 ISC2 Databases and Migration Planning

2.2.1.3.3.1 Current Major Databases

The current major databases are those in the CMU system. Functional applications other than those that were explicitly part of the CMU development generally use data from the CMU databases, and these databases are generally wholly contained within the functional subsystems (Missile Warning, Space Surveillance, etc.) which were federated into CMU when five separate development programs, with individual and separate contractors, were amalgamated into CMU in the mid-1980s.

Table 2-1 lists the major databases included in the current CMU system.

Table 2-1. *Major Databases included in Current CMU System*

Subsystem and Databases	Content	Approximate Size	DBMS	Location
Missile Warning	Missile Launch & Impact Info Missile Parametric Information	250 MB	Flat Files	CMOC and Offutt AFB
Space Mission • Catalogue Maintenance • Space Defense • Static Files • System Control	Satellite Observations & Elements • Maneuverability, etc. • Sensors, Control (relatively static)	4 GB	Adabase	CMOC

Space Defense Server	Satellite elements for web-based distribution to users	1 GB	Oracle	Peterson AFB
Air Mission	Geographic info, Units, Force Status, Targets, Messages	450 MB	Flat Files	CMOC

There have been a large number of problems encountered with the CMU (and other) systems due to the federation of systems with self-contained databases. A sampling and the resource impacts of each are shown in table 2-2 (experience teaches that the actual resource impacts are probably considerably larger than these estimates provided by AFSPACECOM personnel):

Table 2-2. *Examples of Multiple-System Data Use in the N/UWSS*

Title	Discussion	Annual Resource Avoidance
NAVSPACE – AFSPC Space Catalog Exchange	NAVSPACE performs as the backup to the AFSPACE Space Defense Operations Center (SPADOC) in support of the USSPACECOM Space Surveillance and Catalog maintenance Mission. The two systems maintain information in different forms in their respective systems. Because the data is in differing forms, the two systems calculate the catalog independently based on inputs from the same space sensors. With the current catalog size (~10,000 entries), there is an average of 50-100 conditions per day where the external observations received from NAVSPACE require manual intervention to resolve in SPADOC. The catalog will grow over the next 6 years to over 100,000 entries. Extrapolating out, and assuming an optimistic case of only a 2 to 1 growth in the number of manual interventions required for a corresponding 10-fold increase in catalog size, we could expect an additional 100 to 200 manual interventions per day with the larger catalog, resulting in an estimated workload increase of 1.5 man-years.	1.5 MY
N/UWSS - USSTRATCOM Common Information	There is common information that will be shared between the N/UWSS and USSTRATCOM systems in support of the N/UWSS Missile Warning and the USSTRATCOM Force Management/Force Survival missions. Common information includes Friendly Order of Battle data (asset locations and asset categories). These data will each come from a single authoritative source to both systems. Because of the need for consistent results for threat analysis and force survival management, the data used by the two systems must be correct, current, and consistent. An estimated one man-year of effort would be required to maintain duplicate copies of the data, thought this does not include costs of testing with updated data, since that will occur whether the data is maintained in common or separately.	1 MY
Missile Warning – Missile Defense Common Information	There is common information that will be shared between the Ground-based Mid-course Missile Defense (GMD) and ISC2 systems in support of the Missile Defense and Missile Warning missions, respectively. The information includes Missile Order of Battle (missile types, and locations, and associated parametric data) and threat determination data (Geopolitical boundaries, asset locations and asset categories). These data will each come from a single authoritative source to both systems. Because of the need for consistent results of threat analysis, the data used by the two systems must be correct, current, and consistent. An estimated two man-years of effort would be required to maintain duplicate copies of the data, though that does not include costs of testing with updated data, because that will occur whether the data is maintained in common or separate conditions.	2 MY

Title	Discussion	Annual Resource Avoidance
ISC2-TBMCS Common Air Mission Information	Common information required for operations in support of Homeland Defense and the NORAD Air Sovereignty mission will be in both ISC2 and TBMCS. Common mission and mission support information maintained by both is site/facility status of airbases and sensors, mission planning information (Air Tasking Order) and mission execution status. The estimated cost of 1.5 man-years is associated with the double entry and verification of the information in ISC2, and is based on the effort required to enter the same information in TBMCS for a theater approximately the size of NORAD.	1.5 MY
ISC2 (SBMCS) to Other Users – Common Space Information	SBMCS provides common space information to users worldwide in two forms – web-based displays to the user based on user queries and space asset overflight tracks on the Global Command and Control System Common Operating Picture (GCCS COP). If there were no SBMCS capability, then AOC users in theaters worldwide, as well as others, would have to call the Space Operations Center to gain information from SPADOC as a manual task, and plot any data of interest manually on their respective systems in near-real time, resulting in an estimated workload increase of 26 man-years, which includes the impacts to users worldwide.	2 MY for SPOC 24+ MY for Users (assuming 2 MY avg per user and 12 users)
ISC2 Internal Systems Data Management	The current systems (Granite Sentry, CCPDS-R, and SPADOC) use disparate databases based on different commercial solutions (Sybase, Oracle, and Adabase) for their persistent data. Because of the unique features of each and the skills required to support them, they require separate database administrators to manage them. The ISC2-planned integrated solution will provide a single Enterprise Database solution for the three mission domains, thereby allowing a reduction in the number of database administrators required once the system stabilizes, with an estimated savings of two man-years. Common information between the systems includes friendly order of battle (Asset identification and locations). If there is not an integrated solution for Air, Space, and Missile Warning data, then the users will be required to view the different pictures (Air, Apace, and Missile) on separate displays, as is done today, to gain a full understanding of the total battlespace. While the use of separate monitors has been demonstrated somewhat successfully in relatively small command centers, cost and space limitations prohibit application of that approach on an enterprise scale (50+ users).	2 MY

2.2.1.3.3.2 Database Migration Planning

The “To Be” Operational Architecture for the N/UWSS (an evolving document developed by the users) specifies the kind of capability required in the future. The System Architecture, which is developed by the contractor in response to that Operational Architecture, currently describes a prospective system that provides a solution to the current operational requirements based on current technologies. The actual solution will evolve as operational requirements and enabling technologies evolve throughout the evolution of the N/UWSS system. This prospective architecture serves as an evolving baseline design solution, which can then be used as a decision aid to support the implementation of a consistent design solution that takes changes in requirements, interfacing systems, or enabling technologies into account.

Lockheed Martin, in cooperation with ESC, is exploring improved data interoperability techniques based on the Joint Battlespace Infosphere concepts, as currently being explored by ESC and others. ISC2 is using Extended Meta Language (XML) as a standards-compliant, platform-independent exchange mechanism. Developers are building on the broker approach being pursued by ESC and MITRE to provide a sustainable interface. Current plans envision the maintenance of data in databases which are accessed by many functional elements of the N/UWSS system, avoiding the problem that CMU encountered when a set of individually evolving development programs were federated into CMU.

According to the contractor, the Target System Architecture is to be a product line-oriented and standards-based architecture built on proven E-commerce tools. It will provide standard data access mechanisms for all ISC2 data, as maintained in the Enterprise Database (EDB). It includes an Enterprise Workstation (EWS) that provides a common Human-Machine Interface (HMI) for all NORAD, USSPACECOM C2 Missions, and AFSPC, as well as a Virtual Command Center capability that supports access to mission data and information from virtually any location, given proper connectivity and need-to-know. The TSA features what is to be a robust Information Pipeline, implemented through high-performance internal and external networks (The networks are to be government-furnished equipment [GFE]). The internal network architecture plans to leverage the evolving Public Key Infrastructure (PKI) to provide high-assurance data separation on a single backbone, reducing overall network costs and simplifying network management. The external communications architecture planned for ISC2 is common-carrier based, with the objective of providing bandwidth on demand and quality of service assurances to the NORAD, USSPACECOM, and AFSPC missions while reducing overall communications costs by ending reliance on the current costly dedicated point-to-point circuits. Development of the enterprise infrastructure needed to support the first spiral is now underway (the Air Defense function, with AME replacing Granite Sentry, is the operational capability planned for the first spiral). The objective database structure is shown in figs 2-2 and 2-3 (see next page, and page 36).

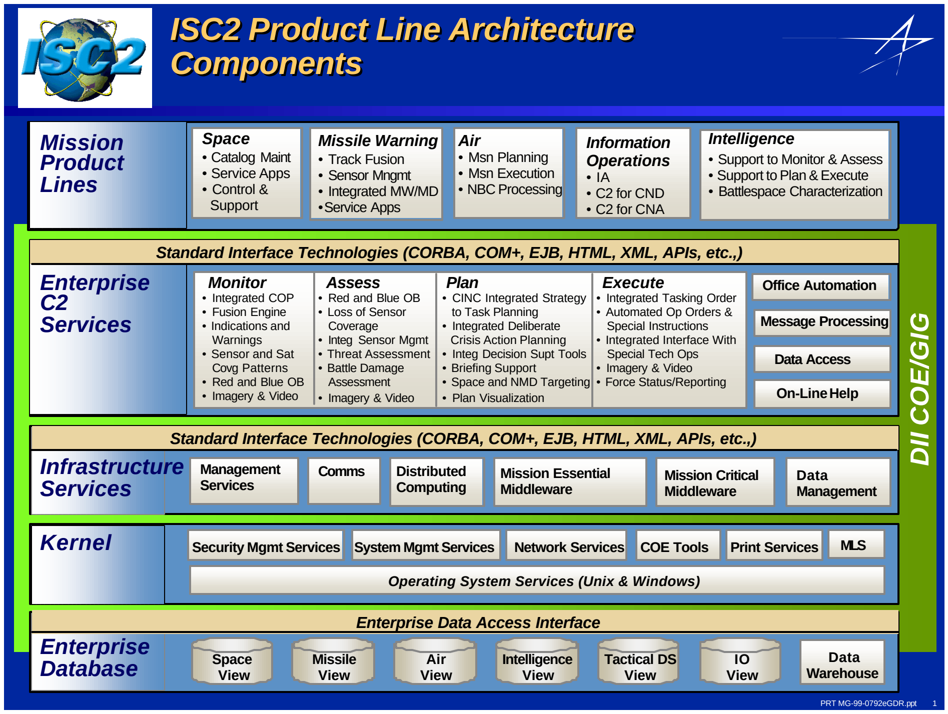

Figure 2-2. *ISC Product Line Architecture Components*

2.2.1.4 Findings on Database Migration in ISC2

- AFSPACECOM elected to define and evolve a whole new architecture for the N/UWSS instead of maintaining and evolving the existing one. A key consideration was the increasing sustainment cost for the existing system.
- AFSPACECOM elected to institute a new acquisition paradigm instead of continuing to use the process that had been used on many generations of legacy systems. The main elements:
  - A program office, staffed mostly from Air Force Materiel Command (AFMC), at AFSPACECOM HQ.
  - Use of task orders and spiral development instead of the "big bang" approach described in the DoD 5000 series directives.
  - Long-term commitment to a single contractor to evolve the system.
  - Use of a "living" Operational Architecture, maintained by the operator, which drives an evolving System Architecture maintained by the Contractor.
- Duplication of data in federated systems of the N/UWSS, prior to implementation of ISC2, requires significant resources.
- The ISC2 Contractor has elected to use an Enterprise Data Management structure, using commercial data management approaches and products.
- The acquisition is still in its early stages and the acquisition methodology remains to be proven (We note, however, that a similar approach has been used successfully for more than five years in the Global Command and Control System [GCCS], with a similar level of funding, about $100M/year. The GCCS program office does not use a Target Operational Architecture, but makes an annual decision on already developed improvements to be integrated into the operational system).

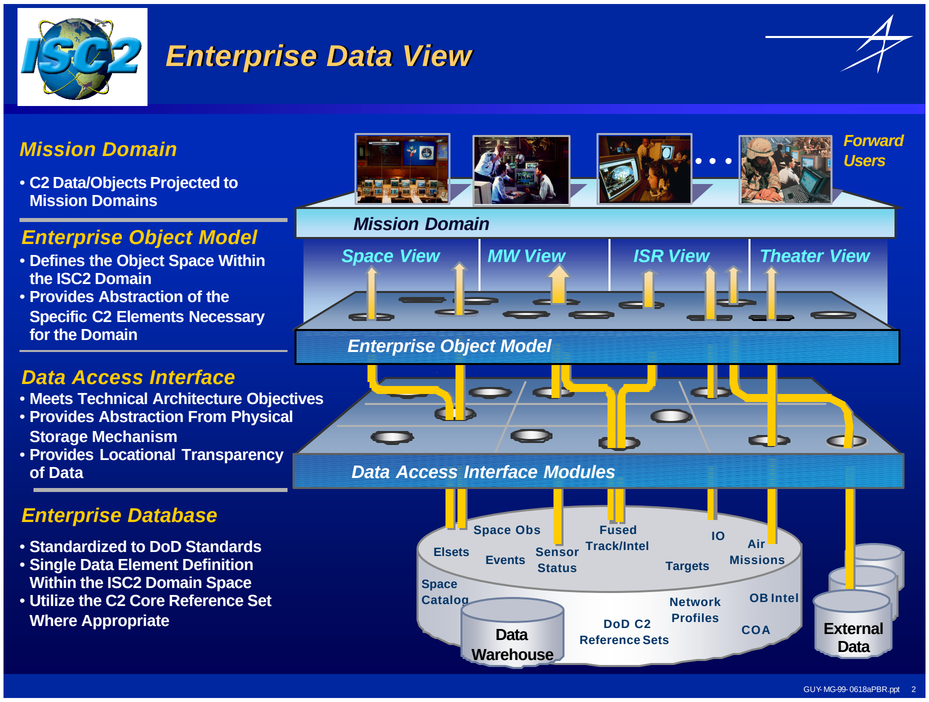

Figure 2-3. *ISC2 Enterprise Data View*

2.2.2 The Theater Battle Management Core System (TBMCS)

2.2.2.1 TBMCS Introduction

The TBMCS supports the general structure of the Theater Air Control System (TACS), working under the principles of centralized control and decentralized execution. AFI 13-109 describes the functions of the TACS and its senior element, the Air Operations Center (AOC), recognizing that each will be tailored to meet the specific and very complex needs of theater operations.

The TBMCS software is used to plan, develop, and coordinate theater air operations. Its key functions include the tasking, publishing and monitoring of the Air Tasking Order (ATO) and Airspace Control Order (ACO). The TBMCS program integrated these legacy systems: Combat Intelligence System (CIS), the Wing Command and Control System (WCCS) and the Contingency Theater Automated Planning System (CTAPS) with the Air Support Operations Prototype (ASOP) using common databases. TBMCS version 1.0.1 is the system of record and received a positive fielding decision in early 2001. Figure 2-4 (see next page) shows the evolution of these legacy systems in the mid 1990s to create TBMCS.

2.2.2.2 TBMCS Background

The TBMCS is the current Air Force flagship program for automating and integrating the planning and execution of the theater air war. Its five core functions can be defined as:

- Intelligence collection and evaluation.
- Planning.
- Generating and distributing the ATO.
- Unit level scheduling of missions.
- Monitoring execution of the ATO.

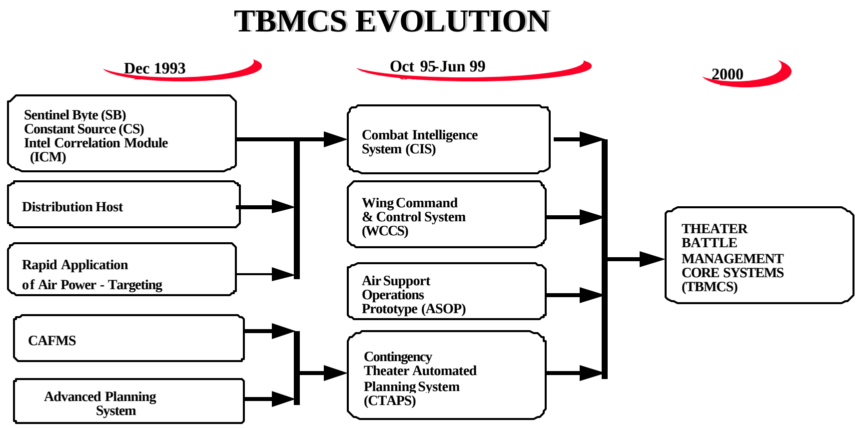

Figure 2-4. *TBMCS Evolution*

TBMCS is intended to link these intelligence, planning, and operations functions through the integration of several legacy systems (or their equivalent functional capabilities), the most important of which are CIS, CTAPS, and WCCS. In addition, TBMCS migrates these key theater air warfare applications to the Defense Information Infrastructure Common Operating Environment (DII COE) platform. The complexity of this integration and migration was underestimated in the mid-1990s when the program was initiated (as have been most, if not all, similar integrations). In the recent words of the then-Program Element Officer (PEO), "It's the most difficult program I have ever encountered."

TBMCS has experienced a troubled and controversial history since its formal launch in late 1995, when Loral (now Lockheed Martin Mission Systems [LMMS] of Colorado Springs) won a six-year competitive cost-plus award fee (CPAF) contract valued around $180 million. The program has suffered from significant schedule slippage, some cost growth, and major performance shortfalls.[3] The original contract envisioned the fielding of three progressively more capable software releases. Instead, as of June 2000, the program still had not been able to successfully complete and field Version 1.0. In addition, the current Version 1.01 represents a significant scaling-down in the capabilities originally envisioned for the first release. As a result, TBMCS is now widely considered to have been a seriously flawed program with regard to its development process, at least in the early phases. The program now, however, generally seems to be on track, and version 1.0.1 has been declared the system of record and is being fielded.

A cooperative effort between ACC and AFMC has established what amounts to a System Integration Laboratory (SIL) at Langley AFB, called the CAOC-X (Combined Air Operations Center-Experimental), to be used in evolving TBMCS by prototyping and experimenting with new capabilities for TBMCS. The CAOC-X was used to configure a TBMCS for the CAOC installed this year at Prince Sultan Air Base, Kingdom of Saudi Arabia — which is currently being used to run the air war in Afghanistan.[4]

Figure 2-5 (see next page) shows TBMCS architecture. It is described as a "plug and play" open architecture designed around a client/server model. It complies with DII COE. There are two primary databases, the intelligence database (Modernized Integrated Database, [MIDB]) and the operations database (Air Operations Database [AODB]), with many lesser databases being used as needed by individual applications. These are

[3] The original contract value to the prime contractor was $35 million (excluding fee, zero base fee), with options that were eventually exercised amounting to $109 million, resulting in a total of $144 million. Award fees and miscellaneous changes raised this to $179 million. A category labeled "evolutionary Requirements (TTDs)" added an additional $161 million, for a total contract value in mid 2000 of $327 million. Mr. Stephen Kent, ESC provided this information.

[4] See AFSAB study TR-00-01, AF Command and Control: The Path Ahead, February 2001 for a complete history of the TBMCS program.

relational databases, with multiple independent and interdependent tables and fields (resembling a three-dimensional environment). The contractor estimated that there are approximately 40 applications and over 500 segments.

Mission applications can be grouped into four general areas. The **Strategic Planning Area** uses the JFACC Planning Tool (JPT) and the Joint Defensive Planner (JDP). JPT provides strategy-to-task analysis and the JDP helps develop and evaluate Theater Air Missile Defense and recommends positions for defensive assets like Patriot missiles. The **Airspace/Intelligence Specialty Team** uses the Intelligence Data Management (IDM) application to interface with the MIDB, and the Imagery Management (IM) application permits access to the imagery servers. The Targeting and Weaponeering Module (TWM) is used to perform target development, targeting, weaponeering, and Battle Damage Assessment (BDA). The weaponeering function uses the Joint Munitions Effectiveness Manual capability to match weapons to specific levels of damage for a given target. **Air Battle Planning** (ABP) uses the Airspace Deconflict (AD) tool, ATO/Airspace Control Order Tool (AAT), and the Theater Air Planner (TAP) to generate the Airspace Control Order and the ATO. The TAP tool supports theater level planning and USMTF messages. **Combat Operations** uses the Execution Manager Control (EMC) to allow field units to input actual data such as takeoff times, sortie durations, and problems with target information, and to follow the execution of the plan. Combat Ops also uses the Execution Management Replannner (EMR) to publish updates to the current and future ATOs. Finally, Combat Ops uses the Close Air Support (CAST) tool to automatically request close air support from either assets on the ground or from CAP.

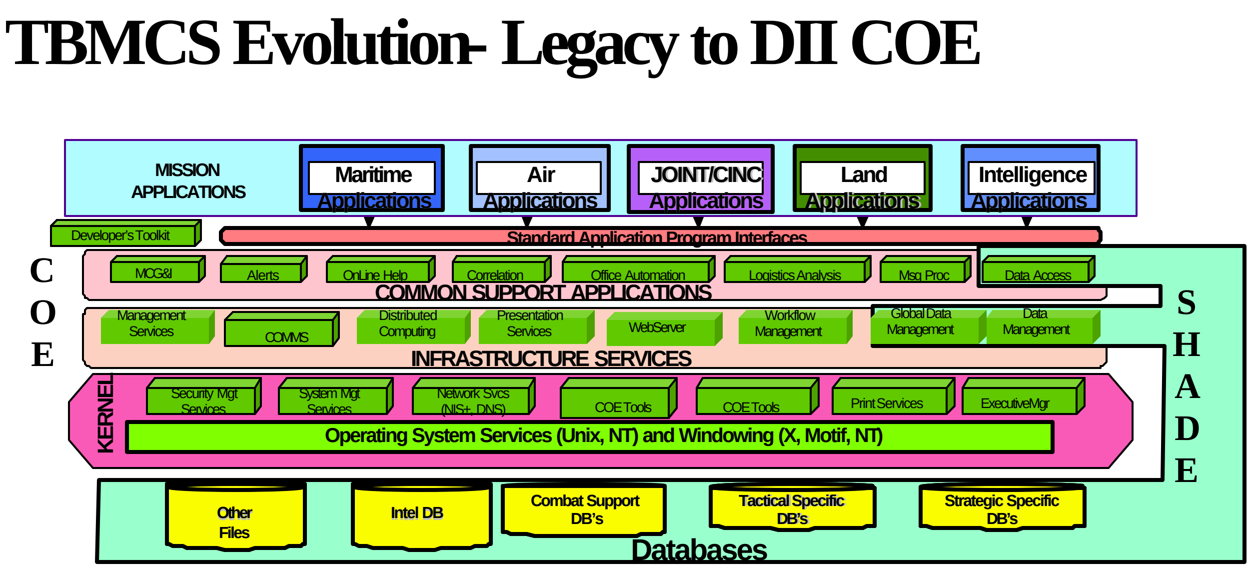

Figure 2-5. *TBMCS Evolution – Legacy to DII COE*

In 2000, the Air Force further distinguished the difference between the unit level and the theater level operation of the TBMCS. PACAF was designated as the developer of the unit-level TBMCS applications. The unit-level TBMCS is not online with common TBMCS databases, MIDB and AODB, but instead copies the databases for their use. PACAF is developing the unit-level software to collect and display the status of the areas of concern to the wing commander. This unit-level portion of TBMCS is, then, basically an evolution of the WCCS.

2.2.2.3 Current Major TBMCS Databases

2.2.2.3.1 General Info on AODB/MIDB

The operator or system support administrator accomplishes needed database management and maintenance. There are four maintenance actions specifically for AODB. For applications, the operator maintains TAP's Friendly Order of Battle and air refueling plan, AIM for airlift, and AD for airspace de-conflictions. The system administrator purges, archives, or restores the ABP. AODB is sized to support a maximum of 10,000 sorties daily for 10 days. The system administrator schedules maintenance down time. Information is accessed by local area networks, DAA services, and query tools such as structured query language (SQL), EMC and embedded SQL.

Additionally, the data management validation rules between databases, specifically for version 1.0.1 and other database providers like the Global Command and Control System- Maritime (GCCS-M) for MIDB, complicate the integration challenge for the program. Enterprise agreement on data management validation would significantly reduce the integration challenge.

Table 2-3. *Combat Plans and Operations Databases*

Combat Plans/Ops Database

Database	Operational Data	Data Model Database Model	Processing (OLTP) & other	Distribution/ Synchronization Model	Database Access
TAP Private	Theater Air C2 Battle Plan	- Oracle 7.3.2 Rdbms - Extensive data model due USMTF 98 migration	- On -line user Upd - Bi -hourly Backup	- Site owns data - Central dB - No data replication - Data synch via ftp	- Corba Tgt, Airspace svc), - Sql - proprietary
JDP Private	**Area Air defense data** **Defended asset list**	**Oracle 7.3.2**	**- Automated process (OLTP) - proprietary** **- Online user Updates** **- Manual backup**	**Pull data**	**- Standard SQL Access** **- Corba airspace svc**
AODB – Production dB	Air Battle Plan Airspace Frob Weather	Oracle 7.3.2	- Auto/Online Upd - Bulk Copy Load/ORDBEX - Hot Backup – Archive login	Procedural (ftp) D istribution	**- Corba Svc** - Std SQL Access - TBMCS svcs
TCT/TCTA	**Time critical target** **Track data** **TADIL warnings**	**Oracle 7.3.2** **Flat File**	**- Discrete transaction with data upds** **- NRT response required**	**Pull mission data**	**Std SQL Access**
Messaging	USMTF messages	Flat File		Role based distribution	Std SQL Access
JPT	Strategy/Objectives	Object Store OO design	On -line user updates Manual Backup	Pull target data	Access Targets via CORBA target service

* BOLD CASE INDICATES CTAPS v TBMCS DIFFEENCES -

Originally, the integrity of the AODB at the various sites would be maintained by auto-replication of the database - a common database design throughout TBMCS. This design allows common application access to common data and provides for improved information availability at all TBMCS sites through the use of data distribution. This is a significant paradigm shift away from reliance on formatted messages for information exchange. This function has been difficult to integrate and has created unacceptable periods of "down time." It has been bypassed by users, who have adopted the makeshift method of manual updates at specific sites. Increment 1.2 had originally been scoped to include use of XML to allow a publish/subscribe capability. Only a few applications have been updated to XML and the update to TBMCS as been deferred.

The incorporation of SQL access and stand-alone databases makes the transparent flow of information and data less reliable. LMMS has decided that the introduction of these tools may allow the user more flexibility, but will make additional data entry a necessity and will further complicate the integration by adding another interface, potentially affecting system performance.

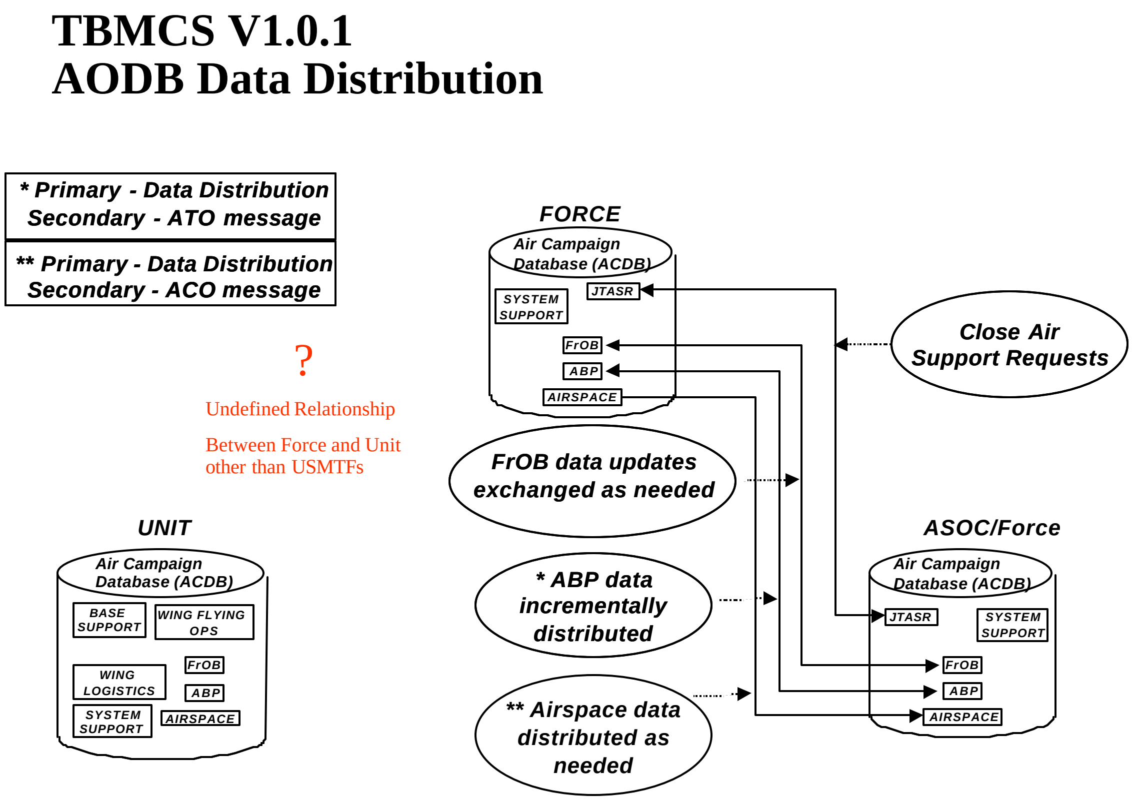

Figure 2-6. *Current Relations Between Databases*

2.2.2.3.2 Intel Database

Figure 2-6 provides a description of the type of data and the manner in which it is accessed in the Modernized Intelligence Database (MIDB). MIDB is supplied by GCCS-M. It is a Sybase database and requires TBMCS 1.01 to use Sybase 11.9.1. Version 1.1 uses Sybase 11.9.2 and Version 1.2 is projected to use Sybase 12.5. There is some work ongoing with development of an Oracle-based MIDB. Use of the Global Command and Control System Integrated Imagery and Intelligence (GCCS-I^3) as a third party acquisition element for intelligence applications resulted in TBMCS fielding older versions of MIDB, which introduced additional interoperability issues

2.2.2.3.3 System Administration

The January assessment team highlighted the difficulties the warfighter has in maintaining and setting up the TBMCS system. This is exacerbated by the Expeditionary Air Force (EAF) concept, the lack of Unix system administrators within the Air Force, the complexity of the network, and the difficulty involved in adding applications. The team identified 40 unique client/server configurations, 2500 profiles/roles, and 300 thick Unix clients with different configurations.

2.2.2.4 TBMCS Database Migration Planning

The government/contractor team is working on an evolutionary approach to continue development of the Theater Level TBMCS system. TBMCS version 1.0.1, currently fielded and the system of record, is not web-enabled. This version of theater-level TBMCS has the common centralized databases and data access agents and is DII COE 3.3-compliant.

The TBMCS program has been restructured to support spiral development and follows the guidance in AFI 63-123. Software Blocks are defined and the current schedule shows Block 20 delivery after JEFX 2002 with the first increment in that Block being Increment 1.2 (described below). The program's efforts to reduce the number of servers required to support TBMCS continue; the number of servers has fallen from 17 to less than 5. The Joint Requirements Oversight Council (JROC) will review the TBMCS Operational Requirements Document (ORD) for comment on 28 June 2001. Future TBMCS increments propose some web-enabled application coexisting with the Unix operating system environment. Some of the web applications are: AAT (ATO/Airspace Control Order [ACO] Tool viewer), Web Scramble, and Web EM Reports. More web-capable tools are planned for future releases but have been deferred from increment 1.2 due to program funding priorities.

Table 2-4. *Intelligence Databases*

Intel Databases

Generic Data Type	**Example Data Types**	**Data Model**	**Processing (OLTP) and other**	**Distribution/ Synchronization Model**	**Database Access**
General Military Intelligence (GMI)	- Facilities, Units, Installations, Targets - Multinode data (e.g., JIC, AOC, Unit)	Relational - RDBMS (Sybase 11.2) - MIDB v2.1.2 from DIA with extensions	- USMTF Messages -Transfer Format: IDBTFs - Auto/Online Upd - Bulk Copy -JIC Load/DEX	- Site owns data - Central dB - No data replication - AOC-unit OB sync - Joint target sync	- Mult. Mechanisms: CORBA-based, (Tgt and OB service), ODBC (messaging), SQL
NRT Track Data Base	NRT tracks	Flat File Structure	- Automated processing (OLTP) - proprietary - On-line user Updates	- COP synch Tool	- Standard API
Imagery IPL/IPA	Images, metadata	Flat File (raster graphics) and relational		FTP uploads HTTP - transfers	Web Access to local store
Weaponeering	Weapons	Flat File	periodic batch updates	None	
Messaging	USMTF messages TADIl A, B, J	Flat File	continuous updates	Role based distribution	SQL
Threat Ring Analysis	Threat Models Track Data OB data	Relational		None	stored procedures

The TBMCS program had a long period of initial operational testing before final fielding of version 1.0.1, due to poor initial results and problems of coordination between the services and the Air Force Operational Test and Evaluation Center (AFOTEC). The initial acceptance testing in 1999-2000 led to an extensive weekly government/contractor team review where the status of open software problem reports was evaluated and risks were assigned. In May of 2000, the program office began briefing TBMCS Modernized Operational Testing and Evaluation (MOT&E) Readiness Reviews. One of the key challenges that the government and contractor conquered was the development of a System Employment Process Guide that would be used for the testing. The testing process for Increments within the proposed Block cycles has not been solidified.

TBMCS version 1.0.2 is currently in testing and uses Sun's J2EE servlets, JAVA applets, Hyper-text Markup Language (HTML) queries and Common Guard Interface (CGI) queries. Version 1.1 is scheduled for Operational Test and Evaluation (OT&E) in October 2001. Increment 1.1 complies with the United States Message Text Format (USMTF) 00, making TBMCS interoperable with NATO systems and reducing the number of servers required to support TBMCS from 27 to 17. LMMS has stated that format changes to the USMTF messages typically drove changes in the AODB database schema and in many applications.

Increment 1.2 incorporates the Joint Targeting Toolkit (JTT) 2.1 developed by AFRL/IF and will be released after JEFX 2002. The incorporation of JTT was directed by the Air Force, but the content of this version is still undergoing refinement. Lockheed stated that they are responsible for integrating the JTT software, but are not responsible for meeting specific performance criteria. Incorporation of XML into Increment 1.2 has been deferred until applications are updated to support XML capabilities. LMMS is not on contract to make those changes.

The Aerospace Command and Control, Intelligence, Surveillance and Reconnaissance Center (AC2ISRC), the program office, research labs, and operational testers continue to make strides toward accomplishing the difficult task of taking advantage of developments from the military and commercial sectors, including off-the-shelf solutions, as well as those successfully prototyped in laboratory or field exercises. The creation of CAOC-X at Langley AFB is still evolving and may provide a formal and cyclical means to quickly integrate new capabilities online. It offers a process for the evolutionary integration of developed modules.

The operators' and developers' main concerns over the past few years have been to get TBMCS operating and established as the system of record. There has been a lot of focus on the disparate database management systems (DBMS) in the system (Sybase and Oracle are both used), and the developers have worked to clear up the difficulties engendered by those disparities. There has also been a major effort to reduce the number of servers required and the consequent footprint of the system. Little effort has been made to reevaluate the data structure for the overall TACS, such as is addressed in AFSPACECOM's ISC2 effort. A draft OA was distributed by the AC2ISRC early in 2001. The understanding that such an OA is necessary to develop an enterprise data structure and manage the evolution to a responsive and reasonably maintainable enterprise system - one that is capable of managing such demanding tasks as the attack of time-critical targets - is still in the future.

2.2.2.5 Findings on Database Migration in TBMCS

- TBMCS was the initial attempt to integrate legacy TACS systems through a "big bang" acquisition approach. Each of the legacy systems had its own database, and there was no operational concept or architecture or even an ORD for the TBMCS system. The operators, developers, and contractors did the best they could. The integration of independently developed systems is very difficult and expensive and is unlikely to yield a satisfactory product.
- The "big bang" approach to C2 systems acquisition is a bankrupt one.
- Further evolution of TBMCS is being attempted using a spiral development approach and a SIL at Langley AFB (ACC HQ). A composite organization to manage evolution is in formation.
- The concept of the AOC as a weapon system (including the evolutionary acquisition of systems such as TBMCS) appears to have been accepted in principle by the corporate AF, but the funding structure and adequate funds are not yet in place.
- There is no apparent attention to data as a corporate asset in the TACS, although the need for such an approach is inherent in the problems observed and the incremental improvements being planned. An Operational Architecture is a necessary precursor for proper enterprise data structures.
- Database migration planning is in its early stages.
- An initial Operational Architecture is being developed (by ACC with ESC support) for the TACS.

2.2.3 Air Mobility Command (AMC) in TRANSCOM's Defense Transportation Enterprise (DTE)

2.2.3.1 AMC C2 Introduction and Background

In the early 1980's AMC (as the Military Airlift Command) began automating its C2 functions to simplify business processes. As these systems matured it became apparent that AMC could further automate and simplify its business processes by consolidating and sharing data between these systems. This was, however, a much more difficult task than it appeared. These automated systems had been developed without regard for the enterprise needs, but with a focus on a select group of users. This approach left AMC with a handful of stove-piped systems that could not communicate and share data.

2.2.3.1.1 AMC C2 Systems as a subset of the TRANSCOM C2

This section discusses the major AMC C2 and transportation systems and how airlift data flows from requirements, planning, scheduling and mission execution to in-transit visibility reporting to the USTRANSCOM Global Transportation Network (GTN). We also discuss the new C2 program – the Global Decision Support System (GDSS) II.

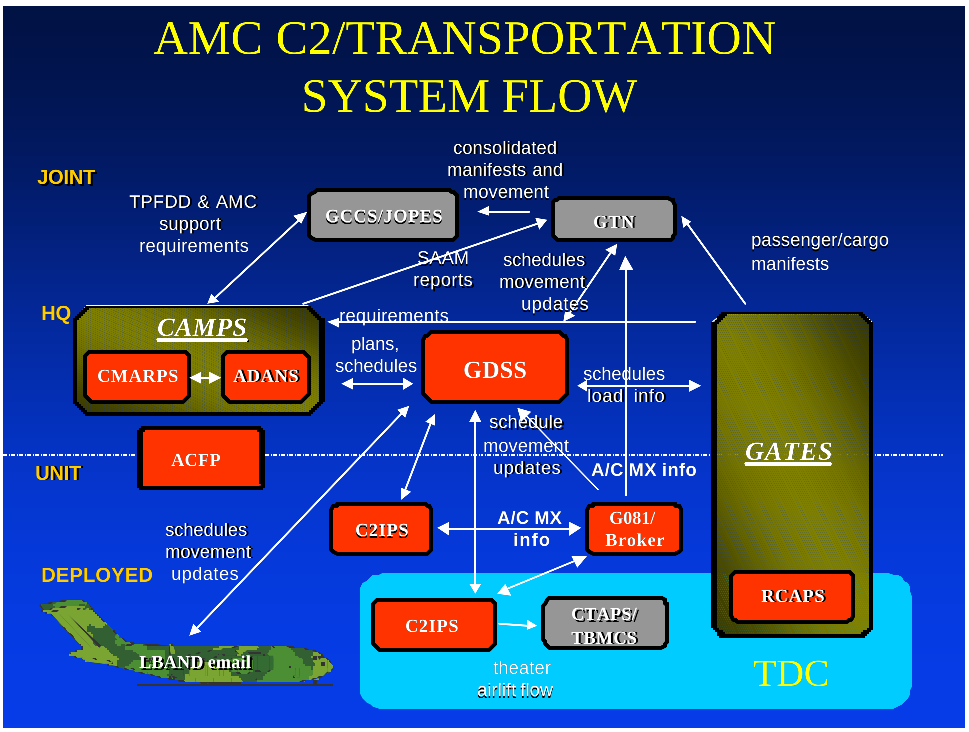

Figure 2-7. *AMC C2/Transportation System Flow*

AMC uses the **GCCS** to receive Time-Phased Force and Deployment Data (TPFDD) information. There are 330 fixed GCCS workstations and 6 deployable workstations at the Air Movement Operations Groups (AMOG) (3 at Travis and 3 at McGuire). Ten of AMC's 12 en route locations have a GCCS capability. Working SIPRNET connectivity exists at the two remaining locations (Incirlik and Rota). Phase IV expansion – Deliberate and Crisis Action Planning Execution System (DCAPES) – will cover the total projected requirement of up to 319 workstations and 227 printers to be fielded AMC-wide in the next 2-5 years. It will also include the replacement of the Contingency Operations/Mobility Planning and Execution System (COMPES) as the logistics, personnel, and manpower-planning tool. This program will affect seven bases and is scheduled for completion by October 2001.

Mission planners in the Tanker Airlift Control Center (TACC) use **AMC Deployment Analysis System (ADANS)** and **Combined Mating and Ranging Planning System (CMARPS)** to plan and schedule airlift and air refueling missions which are then fed into the Global Decision Support System (GDSS). ADANS is located at three sites--Scott, Ramstein and Travis--with over 550 users. CMARPS is used at over 160 tanker and DOD receiver units. These two systems are migrating to the **Consolidated Air Mobility Planning System (CAMPS)**; migration completion date (MCD) is February 2002.

GDSS is AMC's primary force-level C2 system for air mobility assets. The TACC uses GDSS to monitor and manage air mobility and air refueling missions in progress worldwide. GDSS feeds the scheduled missions via the Command and Control Information Processing System (C2IPS) to the flying units so they can assign aircraft and aircrews to the mission. GDSS also passes mission status information to GTN for CINC/joint view of all air mobility missions. GDSS has 6,000 users worldwide. This includes all locations with an AMC presence, ANG and AF Reserve. Primary user is the TACC at Scott AFB. GDSS interfaces with more than 30 major C2 and transportation systems.

Aircrews and TACC planners use the **Advanced Computer Flight Plan (ACFP)** system to generate wind optimized flight plans. ACFP is central to the Mobility 2000 (M2K) initiative improving flight plan filing for central dispatch control.

C2IPS is used to coordinate mission launch and arrival activities during the mission execution process. Its functions include preparation and receipt of airlift schedules and dissemination of taskings. Units and deployed sites use C2IPS to feed the status of the missions back to GDSS. TACC then has a complete force-level view of all missions being flown by AMC and AMC-gained units. The system is also used to support and promote coordination of C2 functions of command post, logistics, and air terminal operations center, and other operations support agencies, primarily through its Sequence of Events (SOE) function. It tracks critical assets (material handling equipment and personnel), identifies and reports all reasons for delayed missions, and tracks and manages aircrew and aircraft resources. C2IPS has undergone a large-scale modernization of its architecture. A new Client-Server (C/S) architecture replaced aging legacy systems. A COTS client (Metaframe) version and Virtual Private Network (VPN) have been added to the architecture, greatly enhancing performance and security. Client-Server is currently installed for 113 Air Force units worldwide. AMC is also responsible for sending strategic airlift mission data to the theater Air Operation Centers (AOCs). The airlift mission data is integrated into the theater Air Tasking Order (ATO), so the Joint Forces Air Component Commander (JFACC) has a total picture of all air resources flying through his airspace. C2IPS maintains a manual one-way interface with the **TBMCS.** C2IPS sends four message types (USMTF format) to TBMCS to provide airlift mission data for the ATO. The current version (3.5.4) of C2IPS achieves version independence from TBMCS. AMC is also developing an automated TBMCS interface and an ATO correlator to allow quicker analysis of the air tasking order for mobility mission information. Both C2IPS and GDSS are Air Operations Center baseline systems and have been installed in the Combined Air Operations Center Experimental (CAOC-X) at Langley AFB.

L-Band SATCOM provides near-real-time C2 of airlift mission execution worldwide. Aircrews use a carry-on laptop capable of messaging directly with the TACC, en route AMC command posts, and regional weather cells worldwide. The ground segment uses the worldwide commercial satellite network, a dedicated circuit to a router/server at Scott AFB. Current e-mail flow is through GDSS C2 messenger for C2 and base e-mail for non-C2 e-mail addresses. Air segment involves permanently installed AERO-C transceiver/printer, antenna, power supply and cabling on 298 AMC aircraft: 114 C-141s, 125 C-5s and 59 KC-10s. Aero-C is an avionics-rated version of International Maritime Satellite (INMARSAT) Standard C equipment already common in the maritime/transportation industry worldwide.

On the transportation side is the **Global Air Transportation Execution System** (**GATES**). GATES is used to enter and track cargo and passenger information. GATES passes cargo and passenger movement data to GTN enhancing total visibility of mobility movement. Cargo and passenger information is sent to the financial systems for billing purposes. GATES replaced the Consolidated Aerial Port System (CAPS II) and several **Remote Consolidated Aerial Port Subsystems** (**RCAPS**). GATES is currently completing the process of replacing RCAPS at 85 fixed locations through FY02.

AMC C2 systems also have an interface with the AMC aircraft logistics system. This provides our C2 systems with the most accurate information on the availability of aircraft to fly missions. **G081** is AMC's version of Core Automated Maintenance System (CAMS). **Broker** is the interface box in front of the mainframe G081 system.

Global Transportation Network (GTN) is a USTRANSCOM system and provides the supported CINCs and other services in-transit visibility into the movement of data for air, sea and land transportation. It is accessed over the NIPRNET and SIPRNET with a standard web browser. GTN feeds associated mission allocation, mission movement, cargo and passenger information to the GCCS Joint Operational Planning and Execution System (JOPES) so that GCCS will reflect Unit Line Number movement and closure. GDSS can also query and receive passenger and cargo information from GTN.

Theater Deployable Communications (TDC) is an Air Force-standard deployable communications and computer service provider. TDC is scaleable and flexible to accommodate mission requirements for mid- to full-scale bare base operations and it supports data rates up to 2.048 Mbps. Sixteen AMC TDC packages will be fielded to AMOGs and tanker wings to support USTRANSCOM and AMC mobility and refueling operations. TDC provides deployed NIPRNET and SIPRNET for GCSS, C2IPS, GTN, GDSS and other serial- or Ethernet-interfaced C4I systems.

GDSS II is a major integration initiative to improve AMC C2 capability by combining the force level functionality of GDSS and the unit-level functionality of C2IPS in a single integrated system, improving mission data integrity and timeliness between force and unit-level echelons and improving reliability and functionality to the user. The goal of GDSS II is a common operational view of air mobility information tailored to the specific needs of system users. A single mobility C2 capability will reduce the need for cross-echelon training, streamline support requirements and reduce program costs.

2.2.3.1.2 Major AMC C2 Databases in the current DTE

The major databases contained in the separately developed but intercommunicating C2 systems described above are as follows:

Table 2-5. *Major AMC Databases*

Database	Operational Data	Data Model	Processing (OLTP) & Other	Distribution / Synchronization Model	Database Access
GATES	In-Transit Visibility Data	Sybase 11.02	User update External interfaces Barcode scanner Auto/manual backup	Replication (Internally & with GTN) Messaging	SQL Web server Stored Procedures
GDSS	MAF Force level execution data	Oracle 7.3	User update External interfaces Auto/manual backup	Messaging Proprietary replication (internally)	SQL Web server Stored Procedures FTP
CAMS FM/GO81	MAF Maintenance data	Oracle 7.3	User Update External interfaces Auto/manual backup	Messaging Brokering	SQL
C2IPS	MAF Unit level scheduling data	Oracle 7.3	User Update Auto/manual backup	Messaging	SQL Stored Procedures
TMDS	MAF Reference Data	Sybase 11.02 InfoPump	User Update Auto/manual backup	Replication SQL	SQL
ADANS	Airlift Planning data	Sybase 11.02	User update External interfaces Auto/manual backup	Messaging FTP	SQL
CMARPS	Air Refueling planning Data	Sybase 11.02	User update External interfaces Auto/manual backup	Messaging	SQL
CAMPS	Airlift & Air Re-	Oracle 7.3	User update	Messaging	SQL

	fueling Data (replaces ADANS & CMARPS)	Sybase 11.02	Auto/manual backup External interfaces		
ABDM	MAF Analysis data	Sybase 11.02	User update External interfaces Auto/manual backup	Messaging	SQL
ACFP	Wind optimized flight planning data	Ingres	User update External interfaces Auto/manual backup	Messaging	SQL FTP
ASIFICS	Airlift financial data	Oracle 7.3.3	User update External interfaces Auto/manual backup	Messaging	Web server

2.2.3.2 AMC Recent Environment

In the mid-1990's, the functional users of the AMC C2 systems complained that C2 messages were being lost between systems and that updates were not received by the Global Decision Support System (GDSS). The communications staff (SC) managers realized that they did not have a way to assess the performance of the data exchange between the major C2 systems. Each Program Manager could produce metrics on how well their system was working, but the metrics stopped at the interface boundaries. The SC managers tasked the Program Integration Office (now the Information Planning Branch) to develop a means of assessing the performance of the data exchange and highlighting areas for improvement.

The Program Integration Office used existing capabilities to establish measurement points and began capturing data on the exchange. The two main areas of analysis were effectiveness (the passage of the message without a receiving system rejection) and efficiency (the time needed for the update to become resident across the C2 environment). In May 1996, the Transaction Analysis team began to publish monthly reports on the causes of message rejections and ways to reduce or eliminate these rejections. The functional managers used these reports and analyses to focus training resources on particular problems. The system managers began making targeted changes in the reference tables and the system validation of entries. The combined efforts began to have a dramatic effect on the message rejection rates. The reject rate dropped from an average of 20% to 50% in 1994 to .02% in 1997. Specifics of how this was accomplished are provided in the Table Management Data System (TMDS) and Command and Control Interface Design Document (C2IDD) details.

As more of the messages were passed successfully, managers turned their focus to the problem of message timeliness. The Transaction Analysis team devised a way to break the entire reporting process into components so they could analyze the components for delays and then address solutions for those delays. The components consist of: 1) the entry process between an event and the user making an entry in the database; 2) the local processing of the message; 3) the transmission of the message through the communications pipeline; and 4) the receipt and application of the update to the target system. The analysis team can compare multiple factors in the process at one or more sites, and can suggest particular areas for further analysis. This analysis can be as simple as an evaluation of the performance of a local network that is either improperly routed or insufficient to accommodate the current workload. Again, managers do not have to use a piecemeal approach to solving performance problems; they can focus their efforts on the root cause to fix it.

Transaction Analysis helped both the functional and the system communities increase the efficiency of their resources. The Transaction Analysis team helps AMC to document the effectiveness of any changes made to the environment. The analysis team produces verifiable results of the effects of a change. AMC was able to quantify the benefits of activating entry validation and the use of standardized reference data. Transaction Analysis helped the program managers secure the funding they needed to improve the performance on the basis of hard evidence, rather than speculation. This continuous feedback loop continues to help functional and program managers set priorities.

2.2.3.2.1 Table Management Distribution System (TMDS)

In 1994, the AMC/SC staff began work on TMDS. The TMDS program had two main goals at its inception. The first goal was to centralize the management of the AMC reference data critical to shared information in the AMC C2 information systems. The second goal was to implement a synchronized distribution of this reference data to these AMC C2 information systems. In the summer of 1995, AMC began distributing reference data in a flat file format for all systems to load at a designated time. The summer of 1996 saw the implementation of an automated distribution process that applied the reference data changes to AMC systems. This automated distribution reduced the time required for distribution of a reference data change from two weeks to less than 30 minutes.

The results of the implementation of TMDS-standardized reference data were dramatic on the C2 metrics. Before implementation of TMDS-standard reference data, the monthly reject rates between C2 systems ranged from 20 % at best to 50 % at worst. After the implementation, the rates fell to a consistent rate of 7 % in 1995.

The keys to the success of TMDS have been centralized management of reference data and the provision of flexible methods for data distribution. In addition, TMDS provides the "24 - 7" support necessary to respond immediately to any crisis requiring reference data changes. As a result, TMDS has moved from its initially successful implementation goal of supporting 40 AMC-specific reference tables and 3 AMC C2 programs to one of supporting 480 reference tables and 15 programs. This includes 180 reference tables for the Defense Transportation Joint Reference Table (DTJRT) project and 6 non-AMC systems.

2.2.3.2.2 Command & Control Interface Design Document (C2IDD)

When AMC (as the Military Airlift Command) began automating its C2 functions in the early 1980's, developers tried to satisfy very limited groups of operational users. This approach resulted in the development of rigidly stove-piped systems using data elements of different lengths, formats and definitions. In the early 1990's, AMC began to document the content and mechanics of the data exchange between the C2 systems. AMC published the first C2IDD in 1992. This first C2IDD documented the exchanges that existed between the systems. In the mid-1990's, AMC began to establish targets at which developers could aim. The results of the C2IDD implementation also had a dramatic impact on C2 metrics. The combination of managed reference data and managed interfaces reduced the message reject rate to an amazing 0.02% in 1997.

The C2IDD now contains the definition, length, format, and business rules associated with each data element that C2 systems share. The Information Planning Branch now holds semiannual technical interchange meetings (TIMs) to determine the best technical solutions to implement the validated functional requirements. These technical solutions are published in the C2IDD about 6 months later for implementation about 6 months after date of publication. This approach gives developers about 1 year to develop and implement the solutions agreed to at the TIMs.

2.2.3.2.3 DoD Data Standardization

AMC's effort to implement data standards is based on DoD directive 8320.1, *DoD Data Administration*, and AFI 33-110, *Data Administration Program*. The desired outcome of the program of implementing data standards is an improvement in data sharing within AMC and with external DoD systems. In fact, DoD directive 8320.1 specifically states in its introduction, "Standard data is the cornerstone of the information infrastructure that supports the war fighter and the overall mission of the Department of Defense (DoD). Sharing information is critical to success on the battlefield and in the supporting functional areas. Standard data will enable DoD to perform its missions in an integrated, effective, and efficient manner."

In 1993, AMC began working on the foundation for implementing DoD data standards. In order to accomplish this goal, AMC began building a logical data model (LDM) based on the data-centric business rules of the organization. To assist in documenting the AMC LDM, a metadata repository was created that was consistent with Defense Data Dictionary System (DDDS). AMC also provided a third normal form (3NF) physical data model (PDM) to aid developers in implementing the data standards.

Initially, successful implementation of these standards was minimal. Developers complained that the standards were inconsistent with the desires of their functional users. The data standards architects complained that the developers were ignoring the standards and that application screens were driving database development. In addition, the products that were necessary to complete an evaluation of standards implementations for a new system were being delivered when the system was ready to field. This left little time for evaluation of a system's implementation of data standards before fielding and no time to correct problems without significant delays in fielding and significant increases in cost to the program.

To address these issues, the AMC data standards architects initiated a prototype partnership in the spring of 2000. The partnership was formed with the system developers for the Integrated Management Tool (IMT) to build a database for the IMT application. The partnership provided an operational database based on the AMC and DoD data standards within six weeks. As a result of this success, AMC transformed its database development process. Now each program is assigned an AMC logical modeler and an AMC physical implementer to assist in the development of the system database. As a part of this new partnership, AMC is currently assisting each program in validating or creating their "As-Is" LDM, creating a "To-Be" LDM and performing an evaluation of the system's current implementation of data standards.

2.2.3.3 The USTRANSCOM Enterprise Architecture

The Defense Transportation System (DTS) Enterprise Architecture (EA) defines the current (As-Is) Corporate Information Technology (IT) Environment that USTRANSCOM operates in, along with the projected (To-Be) IT Environment. The To-Be DTS EA implicitly prescribes the IT environment that ultimately is the USTC Corporate Data Environment (CDE).

2.2.3.3.1 Background:

The "As-Is" DTS EA, published 31 August 1999, represents the current IT environment as a "one-to-many" relationship among systems and data that is application-focused, system-centric and data-based. The "To-Be" DTS EA proposes a target IT environment for USTRANSCOM which includes the establishment of a centralized data management capability and the creation of a CDE in order to achieve an architecture that presents a "one-to-one" relationship among systems that is capability-focused, network-centric, and knowledge-based.

2.2.3.3.2 Discussion:

The global nature of the Transportation Command mission has led to a complex and fragmented environment of command, control, communications, computers, intelligence, surveillance, and reconnaissance (C4ISR) systems and facilities. USTRANSCOM must change the current methods it uses to develop IT solutions to a "knowledge-centric" paradigm. The characteristics of this new environment include global access, and capability-focused, knowledge-based, information-centric, web-enabled and customized application tool sets.

The USTRANSCOM Defense Transportation System Enterprise Architecture (DTS EA) (https://214.3.17.154/dts_ea/) is designed to lay the foundation for a future DTS IT environment and to direct the development of an architecture that enables an end-to-end transportation feasibility capability system. It provides for an integrated, forward-looking, interoperable information systems capability for the DTS community that provides enhanced global mobility, in-transit visibility and competitive rates for the military services and the commercial transportation industry. The DTS EA provides a comprehensive view of the operational processes and C4ISR environment through the operational, systems and technical architectures that document the DTS. The goal is to provide architecture information, guidance and standards to DoD and specifically to the DTS community, to facilitate the integration and interoperability of technologies supporting the DTS.

2.2.3.3.3 Corporate Data Environment (CDE).

The purpose of the CDE is to provide the capability for USTRANSCOM to deploy the tools, resources, leadership, and vision to create and maintain an integrated data platform and easily usable processes and procedures for accessing that platform and for growing it to meet business needs. The CDE is the aggregate of the "infostructure," the set of systems, tools, standards, policies and procedures, and organizations that influence the life cycle of data throughout the USTRANSCOM Enterprise.

2.2.3.3.4 Components of CDE Construction

The components of the CDE construction process are:

- Select and procure an enterprise suite of visualization tools.
- Provide customizable portals to fuse information to facilitate decision-making.
- Move program managers toward applications designed in "N -Tier" Architecture where databases are separate from applications and separate from presentation.
- Build the Corporate Data Solution Layer - Operational Data Store, Data Warehouse, Metadata Repository, Reference Data Repository, and Data Marts.
- Publish and enforce appropriate policy, architectural frameworks and implementation guidance.

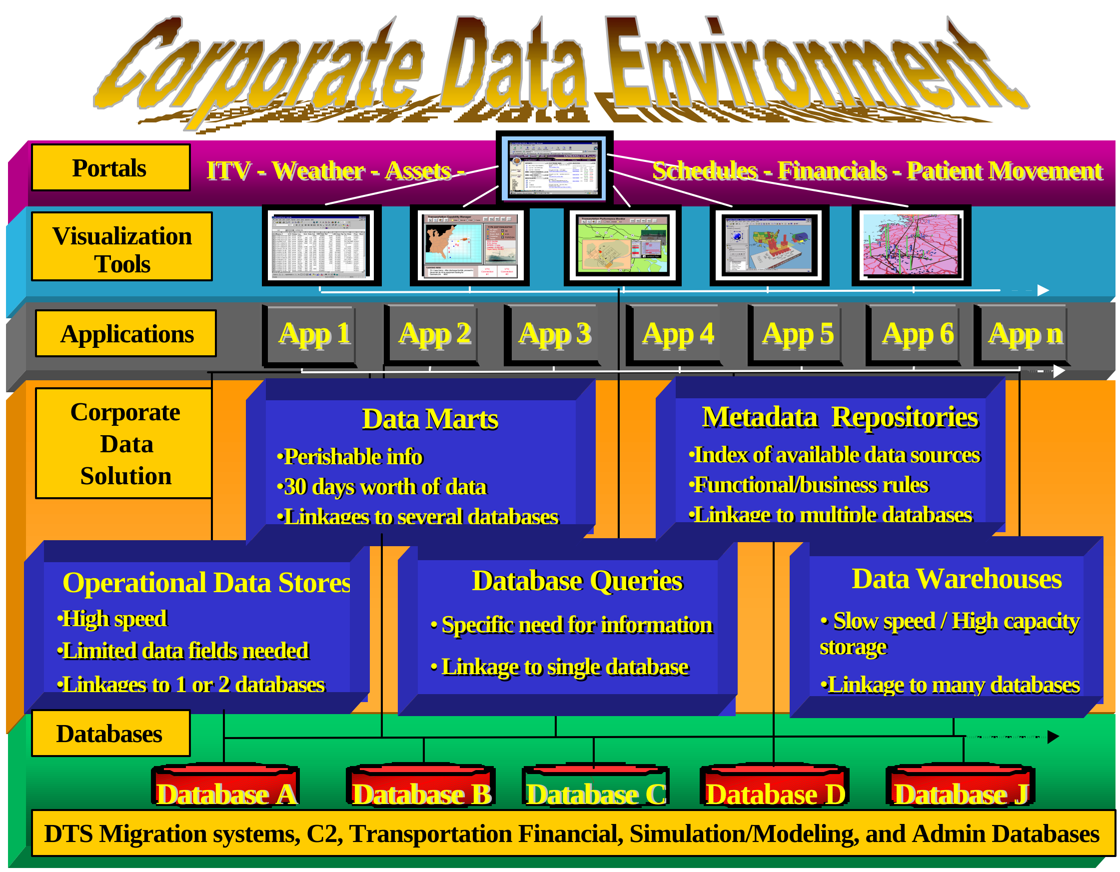

Figure 2-8. *Corporate Data Environment*

The CDE proposes a data-centric management focus, wherein the design, construction, population and maintenance of databases are separated from the development of a myriad of processes that analyze and display that data. Fundamental elements of this structure are:

- A corporate Logical Data Model[5].
- Sound engineering processes.
- An effective monitoring mechanism – agile applications and COTS visualization and presentation tools.

2.2.3.3.4.1 Visualization Tools

A standard tool set is envisioned, bringing together fused information to facilitate decision-making, that users could then call on to customize the desired output for receiving required information from the CDE. Through the use of standard visualization tools, U.S. Transportation Command (USTC) reduces the need for a variety of development efforts and concentrates on the information needed and the format desired to display the information. By concentrating on the information, USTC can forego the overhead cost of the numerous application development efforts. These visualization tools must all meet some common criteria for user acceptance. These common criteria are similar to those demanded by the market place of e-business and result in the best practices in industry. The directed products for the development of all USTC Visualization Tools are XML, JAVA, and Cold Fusion.

2.2.3.3.4.2 Portals

The access to visualization tools comes through customizable portals that provide:

- Single points of entry into DTS databases.
- Access to visualization tools through the Internet.
- Ability for users to customize their view of internal and external data to allow organizations to publish corporate information and allow end users to "subscribe" to the information that interests them.
- Use of intelligent agents to monitor corporate databases for events that automatically trigger emails, pages or voice messages.
- Address the needs of our business partners (Business to Business [B2B] portals), employees (Business to Employee [B2E] portals), and customers (Business to Customer [B2C] portals) through classified and unclassified networks.

Program managers will develop application software to be interoperable with the USTC Enterprise portals. Applications will be developed to provide services via XML for display in the portals. Applications will be developed so that the primary means of access for the user is via the portals. Web-based visualization tools developed by USTRANSCOM, operating on USTRANSCOM networks and accessed through a portal must be built to the following standards:

- Access to multiple applications;
- Application Interfaces;
- Customizable;
- Global Access;
- Information push;
- Single source for Transportation Information.

2.2.3.3.4.3 Applications

Many of the inherent problems in the existing two-tier applications can be overcome by implementing applications with a three-tier architecture. Large, complex projects for which high usage volumes and/or long

[5] See C4ISR Architecture Framework, Version 2.0, Operational View 7

life spans are anticipated will be better served by an N-tier service oriented architecture. N-tier applications have the following advantages:

- Are easily modified to support changes in business rule.
- Are highly scaleable.
- Offer the best performance of any capability.
- Can implement any combination of user interfaces (e.g., character, graphical, web browser, and telephone interfaces).
- Are less expensive to build and maintain because much of the code is pre-built and shared by other applications.

2.2.3.3.4.4 Corporate Data Solution:

The Corporate Data Solution (CDS) is defined as "the components of our infostructure designed to provide aggregated business information." The components are:

- Logical Data Model: representation of data relationships containing entity and attribute definitions and relational business rules (discussed separately).
- Metadata Repository: a database containing the business rules governing the processing of data, the storage structure of data, and the characteristics of data in the CDE.
- Reference Data: data used by more than one system within USTRANSCOM that is not created or modified by the process that employs it (discussed separately).
- Operational Data Store (ODS): a database containing current operational data extracted, transformed and loaded (ETL) from the heterogeneous CDE source systems and integrated into a cohesive corporate view.
- Data Warehouse (DW): a database containing historical data loaded primarily from the ODS.
- Data Marts: a collection of aggregated data summarized from the DW and organized around defined business areas of concern (e.g., location, time, commodity, mode).
- Middleware: software that ensures the application and its data sources communicate quickly, efficiently and effectively, regardless of operating system, communications protocol or database management system being used.

2.2.3.3.4.5 Conclusion

Once realized, the resulting environment of the To-Be DTS EA will greatly facilitate the ability to achieve the CDE. This is, however, a "moving target." New technologies will have an impact on the CDE at USTRANSCOM in several senses. Areas in which new and emerging technologies have been identified include hardware and storage technologies, active data warehousing, data and text mining, agent based technologies, the Extensible Markup Language family of technologies, natural language processing, and the high speed internet. These technologies will provide an opportunity to improve the CDE to better meet the strategic business requirements of USTRANSCOM and its component commands. As time continues, more fidelity will be applied to the ultimate goal of achieving the CDE.

2.2.3.3.5 Reference Data Management in TRANSCOM

The proliferation of unsynchronized and uncoordinated reference tables is a major deterrent to fully integrating automated systems across the DoD. Data errors result when these systems employ different reference tables due to untimely or incomplete table distribution or when table updates are not synchronized across all systems. Complicating this process is the fact that most reference tables are owned and maintained by different organizations, including some commercial entities. Also, many tables are cross-functional in nature. If DoD systems are to be integrated effectively, it is critically important that the processes for the distribution, synchronization and implementation of DoD reference tables be centrally managed by a single focal point.

2.2.3.3.5.1 Background

Unsynchronized data reference tables are a primary cause of significant inter-system data interoperability problems - the scale, scope, impact and complexity of which resemble the Y2K problem. As an example, the GTN alone processes thousands of "bad data transactions" each day that are results of unsynchronized reference tables. An estimated 250 man-years and $13 million are "wasted" annually in processing bad data – just for the GTN transportation mission. A cross-section of other systems reveals that approximately 33% of the transaction errors are the direct result of incompatible reference tables. DoD-wide implications are probably 50-100 times greater than those affecting the transportation community.

USTRANSCOM has developed and implemented reference table management and synchronization processes. The process includes the creation of a "DTS Reference Table Manager" to provide "one-stop shopping" capability for reference table design and content. The Table Manager is responsible for contacting the authoritative source on a periodic basis (depending on frequency of updates) to obtain any changes that have occurred. The manager processes the tables into a relational database and determines the values that have changed. The manager sends the changed values to each of the systems that use the values in the affected reference tables. The process includes policies to enforce DTS Program Management Officer (PMO) adherence to the "standard" set of codes, definitions, and subsequent changes.

On a periodic basis, each of the DTS applications is revalidated against the authoritative source. The application program offices are contacted and asked to send the reference table values to the manager. The manager will compare the application values to the authoritative source. The manager will work directly with the program office to resolve any differences that are found. The key value that the manager adds in the validation process is the ability to work with the program offices in a cooperative manner. The common goal is proactive management of the reference table configuration process across the spectrum of DTS applications.

In 1996 it was recognized that several groups were working on reference table issues in support of Global Transportation Network implementation and on accelerating the electronic payment of transportation bills. There was, however, no single focal point responsible for the coordination of reference table distribution and synchronization for the entire transportation functional area. USTRANSCOM was the appropriate organization to lead this initiative and the Assistant Undersecretary of Defense for Technology Policy (AUSD[TP]) assigned it the responsibility for coordinating the distribution and synchronization of transportation - related reference tables across all transportation systems.

2.2.3.3.5.2 Current Reference Table Efforts

USTRANSCOM is supporting the following reference table efforts:

- Defense Transportation Joint Reference Tables - USTRANSCOM has responsibility for managing and maintaining reference tables across all transportation systems.
- Joint Deployment Reference Tables – "The incorporation of Joint Deployment Reference Tables and systems that support the deployment process into TRANSCOM's reference table management and synchronization process (Joint Staff J-4)."

These programs are funded by the Office of the Assistant Deputy Under Secretary of Defense for Logistics and Materiel Readiness. The pilot program is funded from March 2001 to March 2002. The pilot has responsibility for managing and maintaining logistics reference tables across DoD. The scope of the project is to analyze 100-150 logistics systems.

2.2.3.3.6 The USTRANSCOM Corporate Data Office

During a corporate review of IT investments, the Chief Information Officer (CIO) Program Review Panel (CPRP) in Spring 2000 revealed that multiple USTRANSCOM systems interfaces were created to integrate data for various analytical applications. Some of these interfaces were redundant and resulted in program dollars being spent for the creation of data feeds rather than the provision of enhanced operational capability.

The Corporate Data Environment (CDE), and the Corporate Data Solution (CDS) (its IT solution set) provide architectural principles that will solve this problem. The Corporate Data Office (CDO) was established to ensure effective implementation of that vision.

2.2.3.3.6.1 Background

The mission of the CDO is to establish policies to implement CDS. It will also oversee CDS by synchronizing current development efforts, creating CDS architecture and providing guidance to programs that will build it, and managing command logical data model.

Three major tasks at hand for CDO are: 1) Logical Data Model Institutionalization, 2) Reference Data Management, 3) Metadata Strategy for command. These tasks are listed and discussed in order of priority.

In area of **Logical Data Model Institution**, CDO has two sub-areas:

- Logical Model Synchronization: CDO will maintain Transportation Logical Data Model (TLDM) and ensure that attributes/entities therein are synchronized with DOD Logical Data Model.
- Database Development Consulting: CDO will provide database engineering support to DTS Programs in creating a physical data model that is compliant with TLDM. CDO has built an active partnership with several programs to ensure proper implementation of USTC LDM and will continue these efforts indefinitely until interoperability is achieved.

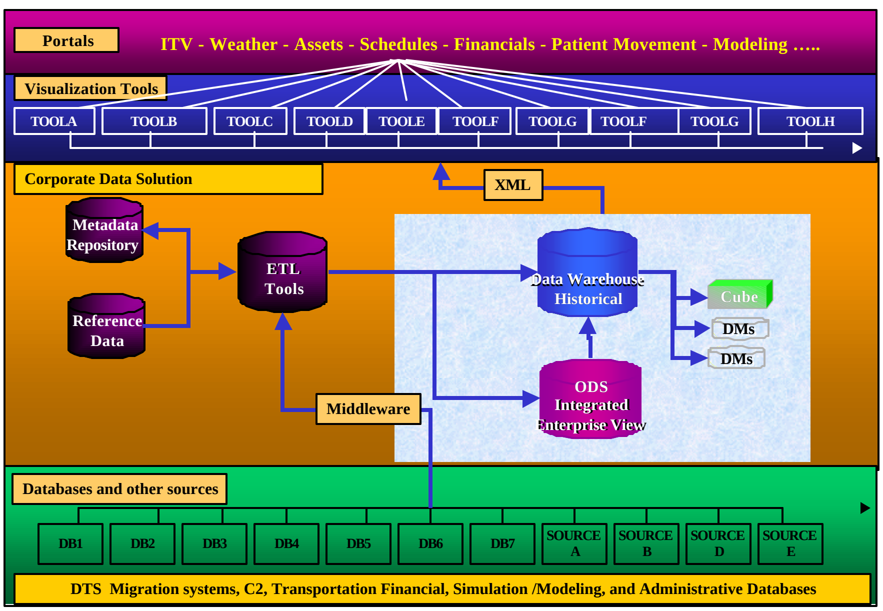

Figure 2-9. *USTRANSCOM Corporate Data Environment*

CDO must be the single point of responsibility for data standardization within USTRANSCOM. As such, USTRANSCOM will gain a single entity for the establishment, maintenance, and enforcement of data standards. To effectively establish CDO with accountability for identifying interoperability issues due to standardization, USTC must use the Technical Review Board (TRB) process. Results of TRB are summarized into a decision-ready package and used to influence resource allocation decisions made by senior leaders at CPRP. In the area of Reference Data Management, CDO is currently manning an office to define and identify reference data. Authoritative sourcing of this data is brokered. CDO works with the owners of TMDS at Air Mobility Command to ensure reference data is properly synchronized across the command.

For Metadata Strategy, CDO will create a strategy that identifies metadata that USTC will need for its implementation of CDS. This strategy will also identify location of required metadata, and a means to create an integrated repository for data. The actual metadata repository will be created, populated, and maintained as a subset of GTN21 and several other programs. The CDO will also define (through the metadata strategy) the process and format for capturing transformation business rules within the appropriate portion of metadata repository. When strictly enforced, an enterprise ETL tool (such as Informatica, DataStage, or DecisionBase Transformer) can be used to bring new data sources into CDS.

2.2.3.3.6.2 Current Integration Efforts

Currently, CDO is in process of synchronizing various efforts:

Global Transportation Network (GTN). The next step in the evolution of the GTN program (GTN 21) will bring about a large part of CDS. With their previous experience with DTS data, GTN will bring about a great capability to extract, transform, and load. Coupled with technology refreshing, the GTN 21 effort will comprise much of the corporate data solution. In the interim, CDO will make metadata available as necessary to command warehouses.

Business Decision Support System (BDSS). BDSS program is currently using Teradata and Informatica products to pilot use of warehousing tools at this command. As GTN 21 nears implementation, an appropriate integration or replacement strategy for BDSS will be developed.

Integrated Movement Data Display (IMDD). IMDD is a system built using the fundamental concepts of CDS. Data requirements for IMDD will become data requirements for final CDS. IMDD serves as an example of CDS success on a small scale. Future data integration projects like IMDD will serve as building blocks of CDS.

2.2.3.3.6.3 Current Projects.

Recently, CIO funded five projects for CDO. These projects are listed in order of command priority:

- **LDM Institutionalization.** Funding was provided for additional data engineers to synchronize Component Command (AMC, MTMC, and MSC) LDMs with USTC TLDM, and consult with programs on implementing standard databases. This USTC TLDM will in turn be synchronized with DoD LDM.
- **Metadata Repository.** CDO will define USTC metadata requirements, identify authoritative sources for this metadata, and implement an integration solution to provide a metadata repository.
- **Data Warehouse Implementation.** CDO will begin an analysis to integrate warehouses across HQ USTC and warehouses of its component commands. By FY02, CDO will have a plan for integrating these sources into a single logical warehouse for the command.
- **XML Implementation**. CDO will define best practices and procedures for using XML within DTS. It will implement a pilot XML project and will provide a full lessons-learned and total cost of ownership analysis to J6-A for the transportation community. A contractor has been contacted for this project and, at writing, should begin work within the next two months.
- **Middleware Implementation**. CDO will define best practices and procedures for using middleware within DTS. It will implement a pilot middleware project and provide feedback to the transportation community of the most appropriate uses for middleware solutions. A contractor has been contacted for this project and, at writing, should begin work within the next two months.

2.2.3.4 AMC Database Migration Planning

2.2.3.4.1 Transaction Analysis

As AMC systems migrate toward compliance with the LDM, the data exchange should get simpler. Transaction Analysis will change its role to one of assessing the quality of data across the enterprise. Their activities will still evaluate the effectiveness and efficiency of data entry and capture, but the focus will incorporate more than just the interfaces. The analysis team will concentrate more efforts on the quality and consistency of the data from the enterprise perspective.

2.2.3.4.2 Table Management Distribution System

TMDS is currently involved with several new development efforts. TMDS is preparing to release a new version that will have improved web access and will introduce more AMC and DoD standardized structures to the TMDS database. TMDS will be adding interfaces to provide reference data to more AMC and USTRANSCOM systems. TMDS is also in the process of adding an additional 450 reference tables to support the Joint Reference Table Logistics Project (JRTLP), a 12-month pilot project for the Defense Logistics Agency. The plan is to use the JRTLP as model for DoD-wide implementation, a move that is currently being considered by the Defense Information Systems Agency (DISA).

2.2.3.4.3 Command & Control Interface Design Document

The IDD will change in two directions. First, it will expand beyond its current concentration on C2 systems. AMC has several automated systems, with more in development, that do not directly engage in C2, but these systems contain information that is critical to the exercise of command and control. These types of systems include functions to document the status of base facilities and services and resource availability. Second, the C2IDD will transition from dictating the format of data elements to specifying the data elements that will be exchanged between systems and documenting the business rules and technical details for that exchange. As systems migrate to LDM-compliant databases, the format and definition of the data elements will become standardized, but AMC will still need a place to document how the systems will exchange or access data, and the business rules associated with that exchange. The IDD will become that single reference source that developers use to understand where data are maintained and how to access those data.

2.2.3.4.4 DoD Data Standardization

AMC will establish a standard workflow for each calendar year for AMC/SC related data standardization activities. Each year will begin with a new release of the AMC LDM. To assist program managers in evaluating the impact, AMC/SC will provide an updated “to-be” LDM for each system based on the newly released LDM. These new “to-be” LDMs will be completed by the end of March. After the systems receive these “to-be” LDMs, they can begin implementing any changes that are deemed necessary to comply with the new LDM. AMC will provide a physical implementer to assist in this process. Another activity that AMC/SC will perform at the beginning of the calendar year is the evaluation of each system’s implementation of the previous year’s AMC LDM. Each system will receive a score by the end of March. The last major activity of the year begins in November. AMC/SC will then collect all the changes that have been identified in that year’s work and begin including these changes for the next year’s LDM. This includes gathering the models and any other required documentation for incorporation into the AMC LDM and DDDS.

2.2.3.4.5 Enterprise Business Rules Repository

One of the latest efforts in AMC is the projected construction of a business rules repository. AMC has captured many business rules. The problem is that these rules are captured in many places. Some of these rules are documented in an enterprise manner in structures such as the C2IDD or the AMC LDM. Many more are, however, captured in system-specific documentation if they are documented at all (other than in application code). This situation leads to inconsistent implementation of business rules.

The Information Planning Branch will develop the AMC Business Rules Repository based upon the standards set forth by the Business Rules Group. This group identified three basic business rule types. The first type of business rule deals with the structural assertions that are equivalent to the entities and relationships captured in the LDM. The second type of business rule is the action assertion. Action assertions are statements of constraint or condition that limit or control the action of the enterprise. The third type of business rule is the derivation. Derivations are statements of knowledge that are derived from other knowledge in the business. Some of the action assertions and derivations are documented in the C2IDD but most are not. The remaining action assertions and derivations are the most critical business rules to be captured in the AMC Business Rules Repository.

2.2.3.4.6 Conclusion

Although each of these individual initiatives would benefit the organization individually, the combination of all four accelerated the successful migration from a group of disjointed, independent systems to a family of systems using common definitions and synchronized procedures. The process starts with the metrics task of quantifying the performance of the systems and highlighting areas for improvement. The maintenance and distribution of standardized reference data helps the systems use the same values while fostering synchronization. The standardized data modeling assists developers from the enterprise perspective by building a common model that defines the relationships between the data. The data standardization provides the formats and domains for those data elements across the Department of Defense. Finally, the metrics task verifies and quantifies the results of the changes.

2.2.3.5 GTN-21—The TRANSCOM Database Migration Planning Focus

2.2.3.5.1 Background

2.2.3.5.1.1 Current System Limitations

In 1999, it was determined that GTN's architecture and design are outdated and incapable of the expansion necessary to meet future DTS requirements. Limitations of the current system include proprietary database design, tightly coupled components, technology obsolescence, complex interface development, unfulfilled requirements, and insupportable architecture. GTN's current database is a COTS product. This product is proprietary. Components of the system are so tightly coupled that even minor changes to one component affect all other components. Technology is moving away from the key assumptions of the current GTN design. Netscape, for example, has moved away from support of the Web Application Interface (WAI) component of its product. WAI is an integral component of the GTN web architecture. Source system interfaces are hard-coded to perform precise tasks. Changes to source system interfaces require a great deal of code analysis, due to their tight coupling with other components. This entails substantial rewriting and regression testing. Due to the complexity of information loading processes, the expansion of the database design, and the system's inability to effectively query the database, many ORD requirements cannot be met with the current GTN system. The current GTN system lacks documentation (the documentation that does exist is poor, there is no LDM, queries are complex, and answers are inconsistent) to support software changes in a cost-effective and timely manner.

2.2.3.5.1.2 GTN 21 Objectives

GTN 21 will acquire a new system that will insert a technology refresher to meet the DTS requirements. GTN 21 objectives have been identified to minimize risks of current system limitations while accommodating data integration issues. These objectives are: the replacement of legacy GTN with modular architecture and current technology, the improvement of current ITV capabilities, expansion of C2 decision support, the minimization of O&M costs, and the ability to adapt to innovative technologies.

2.2.3.5.1.3 Continuing Challenges of Data Integrity for the DTS

Even with the acquisition of GTN 21, there are data integrity issues that will challenge data aggregation within the DTS.

2.2.3.5.1.4 Inconsistent Data Element Definitions

Inconsistent data element definitions throughout the DTS will continue to be a problem for GTN 21. A DISA XML registry, whereby each of 5 organizations uses a different XML tag to describe the same entity, exacerbates the problem. This highlights the disjointed view of business entities across the DoD.

2.2.3.5.1.5 Undisciplined Domain Value Use

The undisciplined use of domain values prevents the maintenance of "clean" and cohesive data. An example of a domain value input a variety of ways would be: C5A, C5-A, C-5A, and C005A. There are also many examples of a variety of disciplines used to fill a particular data element. GTN, for instance, receives more than 50 different time stamps from its source systems.

2.2.3.5.1.6 Location Paradigm

The location paradigm will continue to challenge any data integration effort for the DTS. The location paradigm is exasperated by the propagation of undisciplined business rule implementations and a lack of integrated and enforced location code identifier standards. Business rules regarding location are disjointed with regard to data stewardship, query parameters, and point/region/area definitions. This functional need for location data is complicated by the use of multiple abbreviations, multiple place names, and multiple disciplines for reporting. Some location types in the DTS that need to be policed include: Department of Defense Activity Address Code (DODAAC), Standard Point Location Code (SPLC), Military Standard Transportation and Movement Procedures (MILSTAMP), Geographic Location Code (GEOLOC), International Civil Aviation Organization locator codes (ICAO), International Air Transport Association codes (IATA), Schedule D, Schedule K, Postal Codes, Clear Text Names, and Country Codes.

2.2.3.5.1.7 Date and Time Paradigm

Date and Time Formats challenge data integration because of the global reporting formats and aggregation of the data. Multiple formats and the lack of well-integrated date/time reporting standards are problems for the DTS. Some of the standards used for date/time reporting include: DoD, ISO 8601, ANSI ASC X12, User Defined, and Database Defaults (e.g., Sybase and Oracle). Data and time data elements are not reported with consistent data exchange requirements. Data and time reports are mandatory, conditional, or optional. Date and time requirements are processed using undisciplined business rules. Assumptions are in some cases made based upon system transaction times instead of real-world business event times, which further exacerbates the problem.

2.2.3.5.2 Recommendations

2.2.3.5.2.1 Use of Standard Reference Data

The mandated use of standard reference data would provide a point for managing DTS-common reference data. An office with appropriate responsibilities, power, and resources should provide reference data services. Such services should identify data stewardship, provide oversight of compliance to agreed upon standards as directed by the DoD, and provide a mechanism to correct reference data issues.

2.2.3.5.2.2 Map XML Registry to DoD Data Standardization Efforts

Mapping of the XML registry to DoD data standards would alleviate the problems of multiple names for the same element and elements with multiple meanings. The required use of DoD logical data models in identification of XML names would be a move in this direction. Furthermore, a requirement to map aliases (user everyday-business terms) to the DDDS would provide a pure functional map of the different DoD disciplines to standardized names.

2.2.3.5.2.3 Enable Automated Data Collection

Automated Identification Technology (AIT) that enables data collection at the point of event execution would significantly enhance DTS data quality. Automated data collection reduces "fat-finger" errors. In instances where manual data collection is required, input interfaces must provide more menus of select items, rather than a continued reliance on "hand-jamming".

2.2.3.5.2.4 Enable Data Dissemination

Data integrity is enhanced by the publish/subscribe paradigm of data exchange. The established business rules and processes govern this paradigm. Publication/Subscription efforts provide a clean automated exchange of data between source systems and their customer systems.

2.2.3.5.3 Conclusion

Technology can play a significant role in resolving data integration issues within the DTS if applied at critical junctures. An effort by the DTS should be made to enhance data aggregation by improving enforcement of data standards and applying AIT capabilities for data collection.

2.2.3.6 Findings on Database Migration in AMC C2 Systems—and The DTE

- AMC and TRANSCOM have recognized that data is an important corporate asset and have implemented a process to treat it as such. Early development of an Operational Architecture by the TRANSCOM J-3 has enabled this effort.
- TRANSCOM has set up a corporate data office to assist and monitor the Components' efforts in establishing and sharing DTE data. The office has control of enough funding to insure its clout. AMC is a strong supporter of the effort.
- In addition to identifying data management improvements for the DTE, TRANSCOM is using the GTN-21 program as a vehicle to implement the CDE.

2.2.4 Recommendations

- Do not use the "big bang" approach in developing C2 systems and enterprises. An evolutionary, relatively level-funded approach is much more appropriate for application of the rapidly changing technology that fuels these capabilities.
- Develop and evolve an Operational Architecture (especially the OV-7) to guide the proper selection and dissemination of corporate data within the enterprise. Proper training of the operators (and others) who are to produce this OA is a necessity.
- Establish a continuing organizational entity at the enterprise level to assist and guide the identification and proper maintenance and promulgation of corporate data. This organization must have enough clout to obtain the implementation of its guidance and the responsibility to monitor organizational compliance. Lower level analog may be needed.
- Establish a budgeting and execution structure that allows for the evolutionary development of the C2 component of operational enterprises.

Chapter 3

Attaining Successful Migration

Slide 23: Migration Success Factors

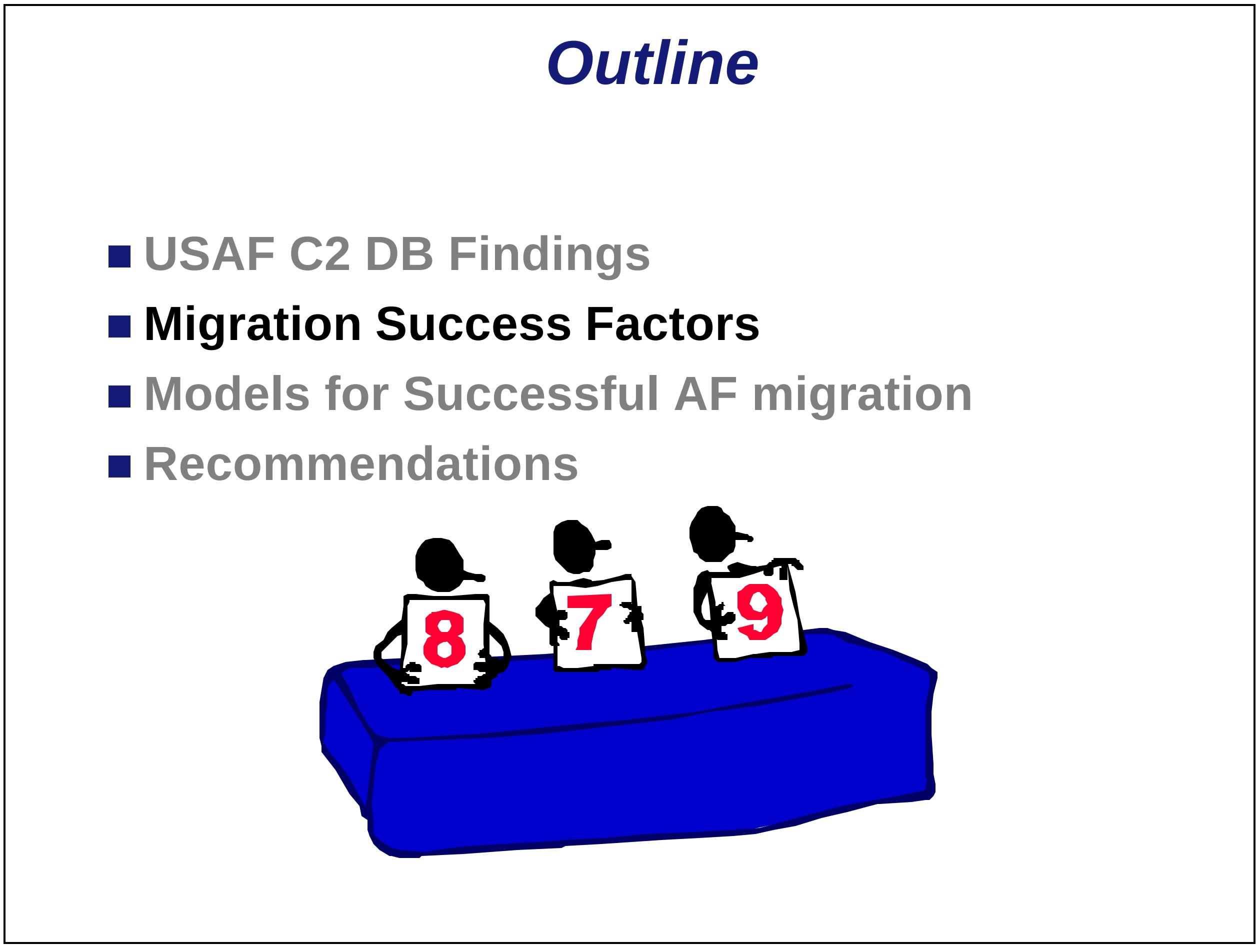

In this section, we discuss some of the corporate “best practices” that we abstracted from many visits and explorations of the current situation in the corporate world. In addition, we surveyed a number of database migrations that failed, and attempted to understand why this had occurred. In virtually every case we found that the successful migrations followed a set of practices that the failed systems hadn’t - and we summarize these in the next few slides. At the end of this section we present a “Technology Summary” that discusses those technologies being used in these efforts.

Slide 24: Background

Background

- Visited institutions with positive and negative experiences
- Visited industrial tool suppliers
- Overviewed new and emerging COTS/GOTS/S&T technologies

Technology	COTS/GOTS today	Emerging
Interoperable data	HTML, DDDS, XML, Weblogic, Websphere	Mediators, Semantic Web, RDF, DAML
Data discovery	Search engines, portals	Information discovery, search agents
Registration	LDAP, Greenpages	UDDI, WSDL
Authentication	Passwords, SecureID, PKI, biometrics	advanced biometrics
Authorization	DB Views, permission bits, classification labels	attribute certificates
Releasability	man-in-loop security guards	automated guards
Legacy Exploitation	Middleware, Corba, JINI, J2EE, ColdFusion	Speakeasy, Open-wings
Resource management	Realtime, Priority Scheduling	QoS, IWAP
Knowledge management	IBM Intelligent Miner, Recon, Lotus Notes	Automated Data fusion
Agent technology	IBM Aglets, Concordia, Odyssey, Voyager, JAFMAS	Agent Grid
System Integration	SAP, Oracle, Siebel, eXcelon	Heterogeneous Access
Integration Components	I2, Tibco, Vitria, BEA, Persistence	XML Services
Database security	Database Views, Trusted Oracle, Firewalls	Role-based access and release control, symmetricsecurity, XML

Appendix C lists all the locations we visited as a group. In addition, our panel included a number of participants who have been involved in both successful and failed migration efforts. We worked hard to explore a number of different technologies, to analyze current COTS/Government off-the-shelf (GOTS) systems, and to explore forthcoming technologies that might hold promise. This table shows many of the areas we explored and the approaches we examined. Although we do not provide an explicit technology roadmap in this report, the table above shows many of the technologies that will help us, if they are properly harnessed, to overcome current database shortfalls. The remainder of this section summarizes the practices used by successful companies in bringing these technologies to bear on their problems. We have particularly focused on those practices that could be easily adopted and used as guiding principles in Air Force C2 database migration efforts. Finally, we aimed at those practices relating to the corporate equivalents of Command and Control, as opposed to issues relating specifically to sales, inventory management, etc. While many of these practices also apply to those systems, we believe the following are the most relevant to overcoming current problems.

Slide 25: Effective Migration Techniques

7 Habits of Highly Successful Migrators

- **Widely adopted package of commercial practices institutionalized by successful companies**
  - **Manage data as an enterprise asset**
  - **Don't migrate monolithically: Evolve**
  - **Determine information flows**
  - **Identify components**
  - **Recognize business-unit competencies**
  - **Standardize look and feel**
  - **Prioritize migration targets**

This slide summarized our findings. We will visit each of these areas in turn on the following slides. At the end of this chapter, we provide a "technical summary" which discusses some of the technical means corporations use to accomplish these best practices.

Slide 26: Managing Data as an Asset

Manage Data as an Enterprise Asset

- **Data must be treated as a strategic asset**
  - **Enterprises depend on decision-quality information**
- **CIO has centralized corporate data responsibility**
  - **Provides enterprise-wide tools and policies**
    - **e.g., enterprise licenses generate considerable cost savings**
  - **Supports and consults with business units**
  - **Trains staff on best practices**

The single most important observation is that the companies we saw recognized that data, and the sharing of data between units, was a critical corporate asset across the enterprise. The only way many large companies can stay competitive is to be able to get decision-quality information to the corporate management at every level (middle managers for their own business units, higher management for the entire enterprise). The first step is the recognition that data from one part of the business can have a major impact on another part. This prompts the consequent realization that a need for data sharing across the entire enterprise exists.

In virtually every case where we saw successful migration in large corporate settings there was a strong management advocate (usually the Chief Information Officer) who was given centralized responsibility for the enterprise-wide corporate data migration and integration efforts. This was especially true in the cases where a merger required the composition of systems from multiple previous organizations, each with different business practices and systems. One of the most powerful tools we found these executives wielded was a corporate enterprise license. The economies of scale make it so that centralizing of licensing can save the corporation a very large amount of money (in the tens of millions of dollars in some cases), and this money can be reinvested for the upgrading and migrating of current systems. The CIO would often not promulgate a policy to make people use these licenses, but would simply supply them “for free” to those who would use them, and would then let those business units who felt they couldn’t use them buy their own licenses for other systems (with the proviso that they interact with the corporate solution). This had two important effects - the first being the creation of a *de facto* corporate standard. Since most units are happy to use the corporate licenses (and thus save their own resources), the bulk of the corporation could be moved to a single approach.

Second, those units that truly needed a different approach (See Slide # 30, “Business Unit Competencies”) were able to depart from an overly restrictive standard as long as they retained the ability to interact with the greater corporate system.

In addition to buying corporate licenses, the central office helps provide training to the business units and staffs on corporate best practices, and provides supporting services to the units which need to help in using approaches and tools different from those they are trained on. Many of the remaining best practices are futile without some sort of centralized control and distributed support and help.

Slide 27: Evolution of Databases

Don't Migrate Monolithically: Evolve

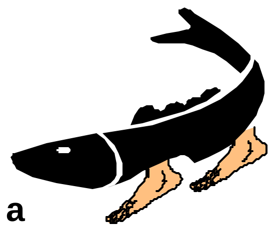

- **Successful businesses use "constant upgrade" approach across the enterprise**
  - **Migrate not to new systems, but to a process of continuous evolution**
    - **Enterprise-wide approach with local execution**
  - **Today's best choice is tomorrow's "albatross"**
    - **Project technological change**
    - **Design against obsolescence**
    - **Wargame against disruptive and competitive technologies**

Since the costs and liabilities of periodic system replacement have become unaffordable in commercial settings of any scale, we find that commercial systems are moving to incremental sustainment. This means that components are upgraded as needed, but well before disasters occur. The overall cost is likely to be the same, but it will be spread evenly over the lifetime of the systems. The initial acquisition will have similar problems. Higher costs may accrue initially, since the needed modularity will require better open interface specifications. "Quickie" solutions to obtain performance are unwise.

A major motivation for an incremental approach is to allow innovations, especially those initiated in the field and hence closely identified with customer needs, to enter the system rapidly. Since such innovations are typically demonstrated in a practical setting, the initial phases of a spiral development cycle have already been achieved. An assessment is needed if such innovations have broader applicability. If they do, scalability becomes an issue.

Some innovations will have applicability to some localities, application suites, and sites, but not to all. The local technological infrastructure, computers and communications may have to be updated to allow for insertion of new innovations. The large costs of global infrastructure updating are spread out over the system's lifetime. Incremental introduction of innovations can proceed as needed. Incremental sustainment means that acquisition will be smaller and more frequent, and the upgrades must maintain the ability to operate with existing software. Program officers will have to coordinate these acquisitions with the managers of extant subsystems. Acceptance testing will become more complex, since there will be more dependencies. Once systems have been tested, however, less subsequent work is needed to build bridges to related systems.

Standards for interoperation pervade such systems and are not only needed for external interfaces. Since standards change, the systems cannot be designed under the assumption that a single standard must serve the entire system. New standards will be introduced as new subsystems are installed, and older standards will be removed. One management responsibility will be to keep the number of standards small, which may mean updating older software even when it is not functionally obsolete in order to get rid of an obsolete standard.

Incremental sustainment is therefore characterized by:

- Incremental replacement of subsystems when appropriate.
- Promulgation of innovations into the enterprise.
- Strategic projects with integrated approaches and explicit support for scalability.
- Coordinated acquisitions:
    - Co-evolving operations and technologies.
    - Permissive rules with respect to standards.
    - Management attention to overall system consistency.

Moving to incremental maintenance of software systems will require a steady and authoritative guiding hand. In substantial commercial enterprises this migration has been achieved, and the ability to cope with incremental requirements and technology adaptations has been achieved. The modules for incremental improvement can come from external or internal sources.

New External Software Modules

The marketplace supports a plethora of software. Selecting the best is difficult. An effort to search for a new software module is only warranted when the functional need arises; it is rare that performance improvements alone warrant a switch to a new software module. The best way to locate and evaluate new modules is through interaction with the community of similar users. "Keeping abreast" is typically a process of participation in specialist conferences, workshops, and the like. We find that the creation of a group that is dedicated to keeping up with external software processes and advances rarely pays off - the required breadth and depth are such that members of such a group must be technically very strong. Such people will not want to work in such a passive role. If no local knowledge exists, consultants can be employed, but it is important to assure their independence.

New Software modules

If no external software covering the tasks is available, then specialized software may have to be written. If the needs existed in parts of the organization then it is likely that some software will already be available.

The latest software tools make it straightforward for novice IT professionals to create adequate software in a manner of weeks, sometimes less. Powerful shareware and freeware tools are a matter of a few clicks and can be built with a variety of add-ons, using no-cost compilers. In the commercial sector, many smaller companies run their entire infrastructure on freeware tools such as Apache, JServ/Tomcat, Perl, PHP, and database software such as MySQL [www.mysql.com]. These tools may not scale well for high-traffic sites where heavy-duty and complex applications servers load-balance thousands of simultaneous requests, but for moderate use they are more than adequate, and may in fact perform as well as their pricey counterparts for small- to medium-sized applications. Even for most military applications, these software tools are "good enough" to handle the simultaneous requests for data, but certainly don't scale when lots of imagery data is requested, and definitely won't work for real-time operations.

Characteristics of Destination Software

It is crucial to consider that the enterprise-level solutions that we must adopt if we are to avoid the problems that stovepiped systems entail will differ from much of the software that is now in operation. Software on the scale that we must consider cannot be created, tested, or maintained monolithically. We must move toward a point where systems can evolve. Evolution requires that components can be upgraded or replaced without making major changes in the system architecture.

Any component, to be replaceable, must be characterized by its interfaces, since those can remain relatively fixed while internal changes are accomplished. To ease replacement those interfaces must be as simple and small as possible and their elements must be well defined. Giving global access to all components, as was done in older architectures and in WMCCS, means that the component must be deeply analyzed before upgrades can occur. If one component requires new database capabilities, then all components must be analyzed before the databases can be upgraded. The demands imposed by changes in such monolithic computer architectures mean that soon no updates at all can be made, and the system becomes an albatross, something that is too heavy to fly and too important to drown.

Evolution also means that we have to exploit legacy software, both the legacy that exists now and the legacy represented by the new software we are building. Nearly everyone realizes that the new software, even before it works, will be legacy code and will require maintenance for repair and upgrade.

Slide 28: Determine Information Flows

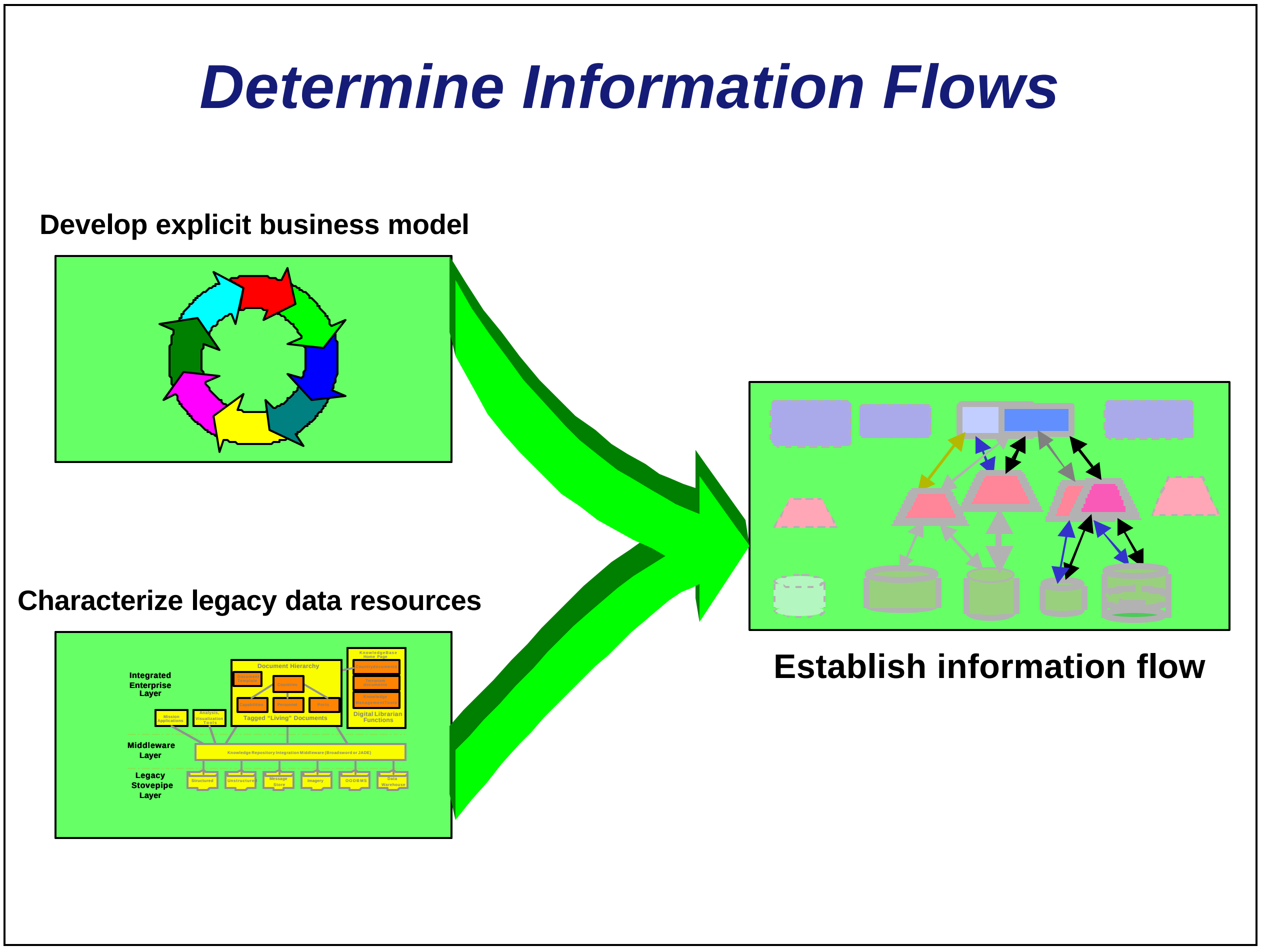

One of the most important tasks that the commercial world addresses when solving this sort of problem (for example, after a corporate merger or when dealing with a major upgrade) is the careful establishment of the information that business participants will need. This task has three major steps - first, the corporation is asked to develop an explicit business model. There are a number of tools available for this effort, ranging from complex modeling tools in languages like UML to much simpler tools used in the construction of a corporate design of operations task differentiation. The second step is the careful characterization of the legacy data sources. Their information flow patterns - the actual information they provide, the recipient of the information and the place where this transfer occurs - are also analyzed. The third step is the identification of information flows. Often, it is determined that some legacy systems are more valuable and better used than others and, in fact, they sometimes find situations where some system is being kept running even though no one actually uses the information it produces.

By performing this task, the corporation is able to determine prioritization targets. It can identify the systems to be turned off first (the least used) and the systems in need of immediate upgrade. More importantly, the information flows can be used to figure out which systems need to be "standardized" and which don't, moving all systems to a common architecture is often prohibitively expensive, and the identification of smaller sets of applications that can be brought together is often sufficient to start the process. In addition, the smaller the set of information systems needing to be made interoperable at any given time, the easier the work of designing an exchange mechanism for them. XML is often a useful tool in such limited interoperation activities; the technical summary at the end of this chapter includes a review of XML technologies.

Slide 29: Identify the Components

Identify Components

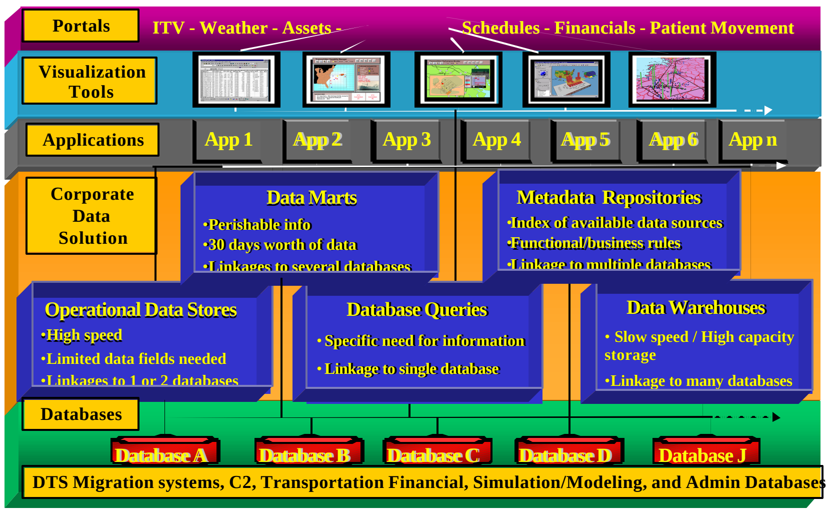

[Ref: USTRANSCOM]

- **Use information flows to determine modular components**
- **When migrating, change components at the lowest level possible**
- **Choose best solution at the right level**

We visited a wide variety of commercial sites during our study. From these visits, we attempted to extract the business practices and habits that made organizations successful. In order to manage the diversity of lessons learned, we placed commercial software efforts into three categories: major system supplier, large niche suppliers, and component suppliers. The names listed in the figures indicate which companies had interactions with us during the study, but this list is far from complete. These suppliers feed each other, even while competing with each other for customer attention. We will summarize their status, as divined by us, in turn.

Major Suppliers.

Major players are those software suppliers who attempt to cover a large fraction of their customer needs, from the underlying databases to the application suites that interact with the customers of the businesses that install these suites. We had multiple contacts with two important players in this arena, SAP and Oracle. These companies each cover a broad spectrum of capabilities. SAP initially focused on turnkey applications, while Oracle started supplying the database infrastructure for applications. Although these companies now overlap and compete for customers, they also depend on each other. For instance, Oracle databases are a primary means of storage technology for SAP. Noteworthy is the fact that these companies themselves, perhaps reluctantly, realize that they cannot cover all needs of their customers themselves. The companies we visited have efforts underway to make it easier for their customers, or the companies that perform system integration (such as Anderson Consulting, now Accenture, or Defense Department contractors such as Lockheed-Martin) to use their systems as part of a global enterprise solution.

Niche players

We categorized as niche players those companies that focus on a particular segment of functions or industries. Their success depends on expertise and the consequent provision of greater added value than the major players can provide. A company we visited in this category was Siebel, a specialist in Customer Relationship Management (CRM). Their success has meant that their definition of customer has broadened considerably, and includes not only outside customers, but also the entities in a corporate supply chain and internal personnel (human) resources.

Niche players that focus on an industry segment have an orthogonal objective, such as the covering of all the needs of a Web-based business (B2C). Such players tend to live "higher in the food chain" and to compete where the major players cannot. Their top-level modules may be portals into a variety of services, but they generate added value by integrating the specialized services that their niche customers require. The domain knowledge they bring to the field often means that they will maintain specialized terminologies, or "ontologies."

Rather than concentrating on technology outside of their realm, niche players will obtain components such as database systems and interoperation technologies provided by other suppliers to the maximum feasible extent. By not developing their own technologies where they cannot add value, they decrease the risk of isolation as technology moves forward. Technology for interoperation is particularly crucial to niche players because they recognize *ab-initio* that they will be but one of the services in an enterprise setting.

Component Suppliers.

The number of component builders is immense. Many component providers rely on specific interoperation standards - CORBA is an example - but have to move rapidly when new standards come about and are accepted by their customers. The customers are rarely the business enterprises that will use the completed information systems. Components are essential to the major suppliers and to the niche players, as well as to the companies that focus on integration services by building customized enterprise systems using niche and component products.

The field of component suppliers is very fluid. A new component supplier may simply spring up or it may be spun off from a larger entity when an internal component appears to have broader value. Some successful component suppliers aspire to vertical growth, and may become niche players. Component suppliers may gain breadth when they believe that they will gain more business by merging with overlapping suppliers rather than competing. Some component suppliers have been purchased by major and niche players to assure supplies or preferential treatment when those players depend greatly on particular component products.

It is at the component level that most innovation takes place. The founding of such companies is motivated by confidence in a new technology or in a new business concept. These small companies will accept risks that are impractical for major players or established niche players. When an established company has an accepted operational process, the uncertainty induced by even considering changes has a high cost. But a new company has no fixed processes, and consequently avoids the costs due to such uncertainty. Once an innovation is proven, the risk of transfer is less.

New components as supplied by innovators have little backup, so some risk remains when new components are placed into larger system contexts. That risk is mitigated by the observation that failures are often associated with decline or non-acceptance of a new technology, so that component replacement is indicated in any case.

Interactions among the layers

Not all companies have clearly identified their role in this "food chain." Many component suppliers may focus on a niche segment as well, and some component suppliers eventually find a niche where their technology has crucial value. For instance, the object-oriented database companies that once expected to support a broad range of applications are now focused on applications in engineering design and software engineering support. Niche suppliers that capture a large fraction of their initial niche will naturally try to broaden their competence into neighboring niches. Major suppliers will initiate new services to compete in enterprises that they previously did not completely cover.

More surprising to us was the extent to which the major suppliers have come to realize that they cannot insist on being the only suppliers for all functions found in major enterprises. Some of them now are willing to provide portals for a variety of services and suppliers, with their own capabilities being but one (although preferably the dominant) choice. As the growth in the field of enterprise support systems continues, new technologies and standards arise and have to be accommodated. It is crucially important that enterprise systems (typically containing parts supplied by all three levels of supplier) retain the ability to evolve. They must also be sufficiently flexible to be able to adopt those technologies that are needed to be competitive. The lesson here for the Air Force is obvious - interoperation and flexibility are the hallmarks of successful enterprises, and commercial suppliers position themselves continuously to fill all the needs and the holes that develop due to technology changes and rising expectations.

Security

We did not receive much information from the commercial sector on the security issues that are so crucial to DoD systems. It does, however, seem clear that while security is a concern, in a commercial setting security concerns cannot be allowed to slow down the dissemination of new versions of the software. One result of this is that commercial systems are not as secure as they could be. Another aspect, and one worthy of emulation if feasible, is the use of simple technologies to achieve adequate protection. Rapid recovery is valued as protection against internally or externally induced failures.

An example where simple solutions are used is seen in banking, where ATM withdrawals are limited in quantity and daily number. Those constraints avoid the complexity and risk of on-line transaction updates for every ATM interaction. The interactions are instead batched and processed at night for efficient and reliable insertion, and the results set the constraints for the next day's operations. At times, but only very rarely, this approach can cause problems for a bank. Such approaches show the benefit of not letting the best possible solution stand in the way of providing adequate services simply and effectively (We discuss security in more detail in the Technical Summary, Chapter 4).

Slide 30: Business-Unit Competencies

Recognize Business-Unit Competencies

- Enterprise knowledge exists at the top level
  - Provides direction and standards
- Specific knowledge exists in the business units
  - Provides operational data and components
- Large enterprises must coordinate global and local competencies

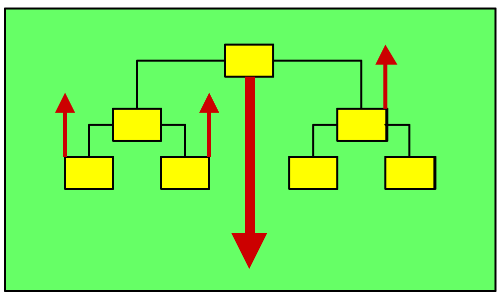

Most big businesses have come to realize what many in the Air Force have known for a long time - the units (business units in the corporate case) often know their jobs better than anyone else. Earlier in this report we cited a couple of efforts where those in the field used common IT tools (such as Microsoft Access™ databases) to create and field prototype systems that better suited their needs than those acquired through the centralized acquisition process. Successful businesses respond to such circumstances by developing processes and policies that allow both centralized standardization and bottom-up development within a set of institutional standards (and the understanding that the standards may be violated, in circumstances of exceptional need).

This is a tricky business, and we did not see any "magic bullet" technology that appeared to work in every case. Management practices that recognized the need to coordinate the global and local competencies were critical to success. The CIO or another central authority generally sets up a process by which the business units and central management were able to communicate at both the technical and management levels. The CIO had the power to decide how to respond to business units' complaints about the necessary tools or processes. One solution was granting him the authority to waive standards when it was clear the business units needed to do something a different way - and when this was done, an effort was made to identify some points of contact between the unit solution and the central licensed architecture. The organization also tried to develop some sort of "shim" which allowed information to flow in both directions. Another solution was the use of a strong interoperability tool or solution that allowed the business unit to interoperate with the central approach - this could require that the business unit do extra work, or a cooperative process (for example to develop a shared XML DTD or schema) might be mandated. Finally, the CIO had the authority to simply overrule the business unit and insist on a particular solution. This authority was used sparingly, but it allowed the expertise of the central group to be brought to bear when it seemed that the business unit was "dragging its feet" for non-technical reasons (such as the maintenance of a favorite system even though the replacement was more powerful and would enable better business practices).

Slide 31: Standardization

Standardize Look and Feel

- **Allow the enterprise to standardize "look and feel," without having to dictate implementation details**
  - **Keep interfaces/control mechanisms as consistent as possible**
  - **Allow customization and specialization instead of over-standardization**

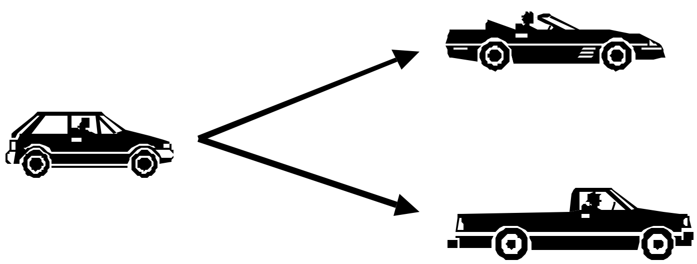

One lesson we learned in the business community was that it, like the military, needs to grapple with the fact that people move from one activity to another. An approach that was used early on was the standardization of systems and practices across business units, a practice that was meant to enable easy interchanges of personnel within the organization. Generally, they found that some units had special needs, requiring waivers of specialized systems that could not easily be blended into the homogeneous whole. This was especially true for those companies that had many of their own legacy systems to deal with, a common aftereffect of corporate mergers. In an effort to overcome this, a number of companies (and the software vendors that supplied them) realized that a better approach was to try to come up with a common set of tools and practices, but not one that required a rigid standardization.

A commonsense analogy is shown in this slide (although it somewhat simplifies the issue). Almost all of the vehicles we drive are standardized with respect to placement of the major control elements - steering wheel, dashboard and pedals, but some vehicles require specialized training. Sports cars, for example, corner differently, and the driver of a light truck learns about cargo placement and special techniques for starting on steep slopes. In a large truck or bus, specialized training is needed to learn about extra gears and more complex shifting rules. In all these cases, however, the basic training on how to drive carries over - so the training burden is drastically reduced, even though the driving tasks differ somewhat. Similarly, software vendors now strive to make more and more systems have an interface that resembles a web browser and the standard "click twice to open" interface seen on personal computers.

It did not appear that the USAF has realized the power this interface “standardization” yields, and many DoD systems still come with custom interfaces that are difficult to master and require far more training than might otherwise be needed. There are some in our community who realize this, although they sometimes seem to call for over standardization - the argument that “all AOCs must look and function alike” is an example of this. Our study suggests that the approach above, of increasing the resemblance between the AOCs’ systems and the operations they support while avoiding the urge to standardize so relentlessly, is the right one. It may be realized economically with a significant improvement in training times and processes without reducing the specific competencies mandated by region or function.

Slide 32: Prioritizing Migration Targets

Prioritize Migration Targets

- **Crucial to prioritize order in which components are to be migrated**
  - **Maintains business viability**
  - **Allows investment flexibility, e.g.:**
    - **Save by freezing legacy functionality**
    - **Savings reinvested later in new functionality**

- **Decisions made as business-case tradeoffs**
  - **Cost vs. risk**
    - **Cost of staying up-to-date vs. risk of missing information**
    - **Assess relative importance to enterprise objectives**
  - **Replacement vs. Upgrade**
    - **Upgrades costly, but not upgrading can be fatal**
    - **Maintenance a dominant cost of DB systems**

The effort devoted to modernizing software systems in the commercial world is immense, and its expense is estimated to be on the order of several hundreds of millions of dollars per year. We were only able to sample a small portion of the commercial efforts, but we made a concerted effort to examine a broad sample of trends. We find that no single technical path dominates, although it is clear that effective operation over the Web is the objective of almost every major initiative. It also appears that due to the lessons learned and the migrations performed on industrial data resources in preparation for Y2K, very few systems of substantial size will be built from the ground up in the future. Industry has learned to view database migration as a constant and ongoing process, not a systems acquisition activity. The value of the knowledge embodied in the legacy systems is just too large, and the risk and disruption caused by major system replacements is unacceptable in a commercial setting. All installed systems must be sustained. Because of commercial pressures, no participant can afford to relax. Keeping systems up-to-date so that customers will remain satisfied is a task of great importance. Without customers, (i.e., coverage and market share), income will drop below levels necessary for corporate survival.

The maintenance costs of information software in the commercial world are at least as high as, and often higher than, those in the military. We repeatedly encountered estimates of lifetime maintenance costs that ran to 70-90% of total software costs, including acquisition and installation. Hardware costs, on the other hand, have been dropping rapidly. Hardware today requires little maintenance, while software costs have decreased little, so that overall costs are dominated by the cost of software and software maintenance. Such a level of maintenance is not alarming, for it represents the cost of retaining a long-term investment that is essential to the enterprise. In some ways it is comparable with the maintenance cost of other long-lived equipment, such as the B-52 or the C-141. Software maintenance costs are mainly induced by the need to keep up-to-date, both with advancing computer and communication technology and with customer expectations. Bug fixing is a minor component of maintenance, once systems are in place, unless unplanned demands are made on older software.

The cost of keeping large information systems up-to-date is high. However, the cost of not staying up-to-date is higher, due to the loss of benefits, inability to access new sources, dependence on obsolete technology, training costs for unpopular interfaces, and increasing error rates. In the commercial world, these losses translate to customer dissatisfaction and loss of income; in the Air Force they translate into the sorts of significant command and control errors described earlier – errors that can lead to loss of life and military failure.

Slide 33: Summary of Best Practices

Best Practices Summary

- Successful migration approaches have been identified and used in industrial practice
  - Manage data as an enterprise asset
  - Don’t migrate monolithically: Evolve
  - Determine information flows
  - Identify components
  - Recognize business-unit competencies
  - Standardize look and feel
  - Prioritize migration targets

USAF should adopt these successful migration practices

It is clear that we must find ways to apply these lessons to the USAF, if we are to gain from the best practices shown us by industry. That said, some of these tasks may be rendered more difficult by the special needs of the DoD, and we discuss these in the next section of the report, following a brief review of some key technologies used in applying these practices.

(This Page Intentionally Left Blank)

Chapter 4

Relevant Technologies

Slide 34: Technology Summary

4.1 Technology relevant to Air Force Information Systems

There are numerous data, information, and knowledge management technologies that are evolving and are relevant to this study. The technologies of interest to the Air Force for modern approaches in information management include those needed for the migration of databases and their applications in a common data environment. Aspects of the systems critical to the Air Force include security and information control, real-time processing, and management of unstructured data. In this section we will first provide an overview of the technology evolution path and then discuss some of the details of the key technologies that will facilitate C2 Database Migration. We essentially synthesize the relevant technologies seen in DoD and commercial practice. The result is a summary of the tools and methods that are available to support the approaches that this report advocates.

4.2 How to Migrate

Migration is the process of moving from one set of software to another, a process often motivated by improvements in hardware or functional requirements (see Appendix A for a discussion of migration). In today's Air Force context, this requires the migration of legacy applications to the shared data environment foreseen in the SAB's JBI studies. The central feature of this modern approach to C4I is the *infosphere*, a global information infrastructure that supplies a fused representation of the battle space in real time. This vision assumes a high degree of data sharing between systems; however, many information systems in the DoD are *data stovepipes,* built to meet the information requirements of their users, with much less concern for the requirements of others. In most cases these systems cannot be discarded and replaced. They must instead be gradually migrated towards a world where data is shared. This section describes a strategy for turning "stovepiped" systems into participants in a particular kind of shared data environment, in which the application programs become clients of shared *data servers*. These data services can be local or remote, so that rapid sharing of information over modern computer networks is enabled, replacing the now common periodic replication of data and the need for redundant entry of the same data into multiple systems.

The strategy outlined is predicated on a move to modern commercially validated processes. Letting multiple applications gain access to shared data is an old concept, and it motivated the WMCCS effort in the late 1960s. Today, however, commercial standards have developed that support the interfaces, so that no dependence needs to be made on proprietary software and hardware. The performance of the hardware has improved so that shortcuts to gain efficiency need not be employed and mediating modules can be inserted in the middle between the applications and the data sources so that flexibility can be gained. In the next section, we sketch the processes that allow legacy applications to share local and remote data resources, and we also address some of the issues that must be dealt with during the implementation of a migration.

Note that the technology described throughout this chapter is intended to address large, shared information systems. Not all data used in the Air Force is heavyweight. There is a place for lightweight, open-source storage systems, where storage demands are modest and encapsulated. These lightweight data systems are omnipresent. Large numbers of DoD users work with small Access™ databases on their machines and these are often needed for the performance of larger missions. We will see in Section 4.3 that mediation can be neatly used to query a mix of heavyweight and lightweight information systems.

4.3 Data sharing and shared data servers

Most information systems create and maintain private data. This data is not shared because it is of no interest to any other system or organization. The opportunity for shared data arises when several systems have an interest in the same concepts. *Shared data* means that the systems use the same data values to record facts of mutual interest. Conceptually speaking, there is one copy of these shared data values; practically speaking, the data may be distributed or replicated over several locations. Figure 4-1 represents the private and shared data of three systems. The shared data is shaded; private data is not. Note that in this picture, shared data is data shared *anywhere*, and not just data shared *everywhere;* it is the union of the intersections, not just the intersection of all.

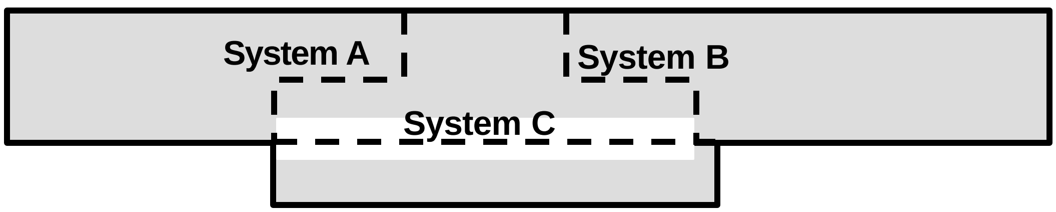

Figure 4-1. *Shared and Private Data*

Stovepiped systems are constructed with separate databases. System developers define separate, independent data schemas for these databases. This means that even where several systems have overlapping interests and could share information, the different names, definitions, and formats in the schemas cause the same *information* to be stored as different *data*. For example, Figure 4-2 shows three data schemas that describe the same information about aircraft maintenance, but which use different names and data structures to do it.

System A			System B			System C		
Table: ACMAINT			Table: RDYACFT			Table: MAINTSCHED		
ACTYPE	RDYWHEN	NUM	MODEL	AVAILTIME	QTY	RDYTIME	F15S	F16S
F15	0500	22	F15	0500	22	0500	22	-
F16	1700	16	F16	1700	16	1700	-	16

Figure 4-2. *Same Information, Three Different Data Schemas*

These three systems have *shareable information,* but cannot share data until some or all of them alter their data schemas and modify their application programs to use the altered schema. As we will see, this is a difficult task. When stovepiped systems must communicate, *ad-hoc* data interface programs are created for each pair of communicating systems. These interface programs convert data from one schema to another as necessary. Often manual operations are needed to synchronize the information, since operational settings and availabilities differ. This situation is illustrated in Figure 4-3 (see next page).

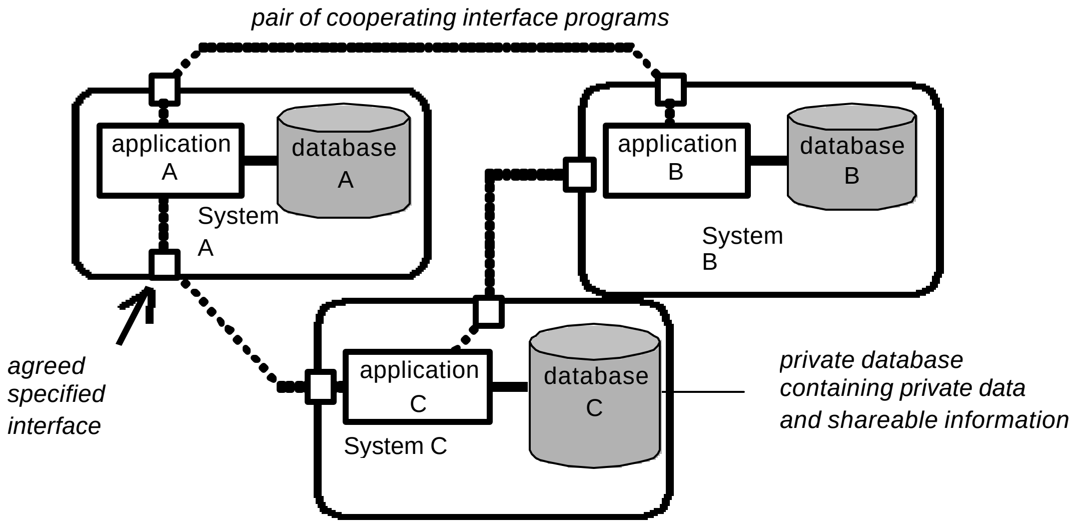

Figure 4-3. *Private Database and Pairwise Data Interfaces*

Experience shows that development and maintenance of these interface programs is expensive in terms of both time and money. Worse, the total effort required increases with the *square* of the number of communicating systems. Finally, these hard-coded interfaces support only the information transfer requirements defined at system development time, and not the "pull-on-demand" transfers anticipated in the Infosphere. After migration, the Air Force systems will communicate through shared data.

4.3.1 Migrating Local Data Systems

For local applications, the best way to share data is through a shared data server. Application programs then become clients of this data server. The formerly distinct data must be integrated. This integration process involves three tasks: developing a data model and data elements for the shared data, converting the legacy data values to this new representation and modifying the application programs to use the shared data server and its schema. During integration, meaningful differences are often discovered. One application may cover a different subset of the information; the granularity required by distinct applications may differ, as may the timeliness and accuracy requirements.

A first step taken by some systems is to move the existing databases to a shared data server without performing any schema integration. A collection of system integration rules known as *segmentation* lets the separate databases reside on a single database management system (DBMS) without mutual interference (Segmentation rules for applications are part of the Defense Information Infrastructure [DII] Common Operating Environment [COE]. Segmentation rules for data are forthcoming). The result is sharing of the data server but no sharing of data. Private databases are simply separated from the client systems and stored as separate databases on the data server. Communication still takes place by means of pairwise data interfaces. Access by remote applications is, however, enabled.

The data interfaces are gradually replaced by queries and updates to the shared data. Each set of applications is allocated a view into the shared data. The views will overlap where information is actually shared. View elements may be restricted to be "read-only" to protect authenticity. Some systems may also retain separate databases for temporary and fully private data. This objective is illustrated in Figure 4-5 (see page 84).

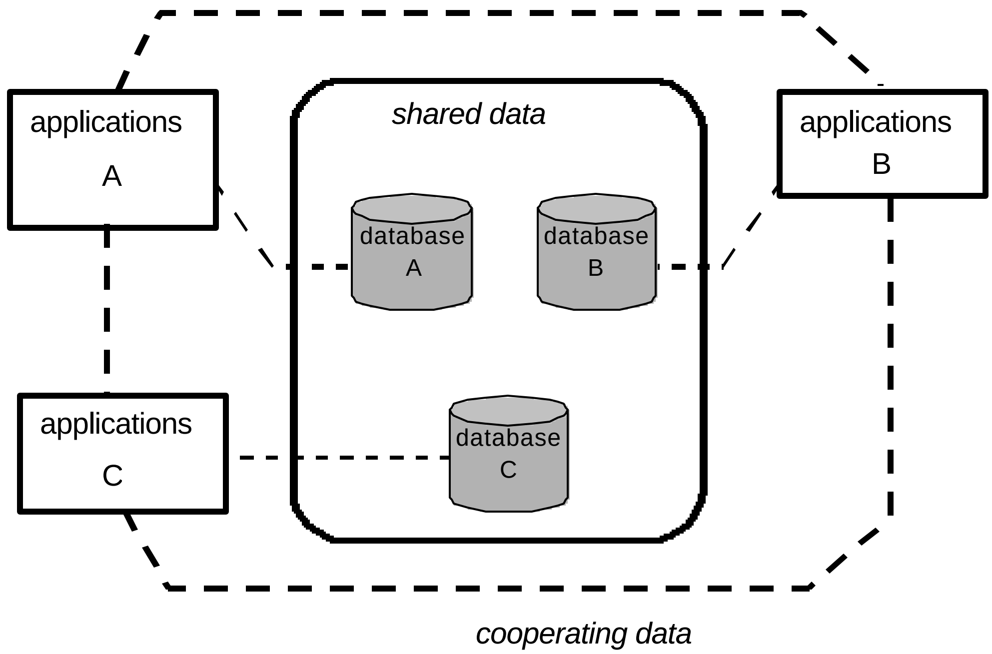

Figure 4-4. *Shared Data Server Without Data Sharing*

Migrating from a system of data stovepipes to one of shared data is difficult. Building the shared, integrated schema is another difficult problem. Few existing applications can be discarded and much valuable knowledge is encapsulated in the code. The existing data cannot be discarded because it represents important historical information. Extended downtime during cutover cannot be permitted because many of these applications are mission-critical. Finally, a single simultaneous cutover of all applications cannot be required. It can be arranged only rarely and the risk of failure in such an "all-or-nothing" migration is unacceptably high. These are problems that must be addressed by any migration strategy.

4.3.2 Mediating Technology

To minimize the impact on existing applications, a data mediator can be inserted between the existing application and the sharable database. Mediators serve as *gateways* that make the legacy data available through the shared data server in terms of the shared schema. A data mediator converts queries and data from one data schema to another. A query from the remote application's schema is translated into the equivalent query against the source schema. Then the source query is executed and the retrieved source data is translated into the receiver's application format. The mediator acts as a *semantic gateway* between the systems, permitting the receiver to view the source as an extension of its own database, without concern for the differences in names and representations of data.

The process of moving the actual data from the original stovepiped systems requires care as well, but it can be done. Historical data should be inspected and cleaned before transfer to avoid the transmission of localized errors from the small stovepiped system to the larger shared system. When data is needed that is not yet in the shared database, a data coordination function in the mediator can issue a subquery to the appropriate source, retrieve the data, and merge it into the result. At least a few weeks of smooth operation should be accomplished before subsequent stovepiped applications are moved into the shared environment. Again, an incremental scheme for maintaining information systems will distribute the cost and loads over a longer period while allowing for a higher quality of the overall operations.

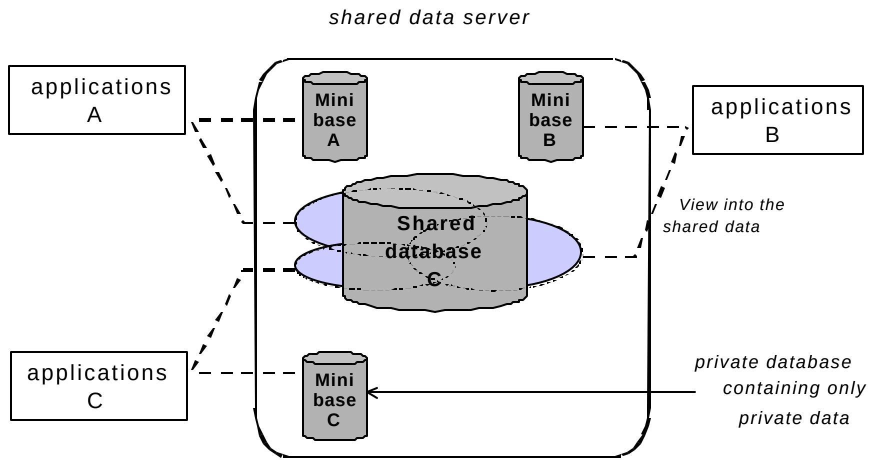

Figure 4-5. *Data Server With Data Sharing*

When, in cases such as the Air Force's C2 DB systems, the sources of information are widely distributed, the integrating solution described in Chapter 2 is no longer feasible. At each source, integrated or not, views can be defined that specify the date to be made available for shared use. Using a data schema for the shareable data, a view of the shared data can provide a mapping for remote applications.

Here data mediators are nearly always inserted between the remote client applications and the shared data servers to allow the combination of information from multiple sources. Since synchronization of remote system migration is rarely feasible, the mediators must deal with legacy as well as with more modern resources. The incremental approach that is now enabled also allows systems to be gradually updated when modern systems eventually become legacy resources.

Mediators can also combine information from multiple local or remote sources as indicated in Figure 4-6 (see page 85). In those cases, the software in a mediator will not only match formats but will also match the granularity of the data in diverse sources, translate data values where they differ and resolve problems due to differing update times in the sources needed by higher-level applications.

Once this has been accomplished, mediating gateways are created for applications that translate their queries and translate the results obtained from their views of the shared database. When read requests overlap in the shared database, redundancy of data is reduced and consistency is improved. When multiple applications update information that is now in the shared database, overlaps must be removed. One application should be given the primary authority and responsibility for updating each data element. The mediator can filter out unwanted updates if the application can be changed to avoid them. In the rare cases where multiple applications must have update authority for a data element, the mediator can become the arbitrator, and can chose the best, the most recent or the most trusted values for the shared database.

Mediators have many useful properties. It is possible to build mediators that can query "lightweight" databases such as the Access databases that are omnipresent in the DoD. Furthermore, mediators can be "stackable" in the sense that some mediators can consider other lower-level mediators as data sources. One may therefore use a heavyweight mediator to query shared data servers, use one or more lightweight mediators to query distributed, lightweight sources, and use a third mediator to query the data accessible through the previous two mediators.

4.3.3 Migrating Global Data Systems

Such mediators can reside anywhere in the network- near the sources, near the applications, or at a site where the expertise exists to deal with that type of shareable information. Since sources and applications change continuously in large information networks such as those envisaged by the Global Infosphere, maintenance responsibility for the mediating nodes must be clearly assigned. It is often beyond the capability of the sources to understand how their data will be used, and beyond a specific application to understand the breadth of relevant information.

4.3.4 Accept Heterogeneity

An important principle here is that heterogeneity must be acknowledged and dealt with, rather than insisting on moving to a level of consistency that is unattainable in a modern and changing world. Application systems can share information that comes from new as well as legacy systems and information sources can include databases, traditional files, so-called "stovepiped" applications and any other computational resources available in today's networked military world. By making the issue of heterogeneity explicit, the problems can be dealt with, rather than be used as excuses for failing to migrate to a modern systems approach.

Heterogeneity comes in many guises. Some are easy to deal with, some are remedied by adherence to standards, but others are intrinsic to the fast face of development of capabilities and needs. We list the principal types of heterogeneity here, with a brief note on how each type may be dealt with.

- **Computer hardware:** Variability is decreasing and common external architecture standards are increasingly well accepted. It is worth noting that heterogeneity in operating systems and platforms is an acknowledged positive in information assurance and prevention of information attack.
- **Computer operating systems:** Variability is decreasing here as well, and common interface standards are increasingly well accepted. Again, however, the increased danger of viruses and other information attacks may militate against a completely homogeneous computing environment.
- **Communication methods:** The Internet infrastructure has made sharing a de facto means of operation, and many commercial suppliers provide methods. There are a number of special case communications methods deployed in the Air Force (and other military) systems, but more and more these are becoming interoperable with commercial approaches.
- **Security requirements:** The requirement that merged data be classified at the highest of the source levels leads to costly operations. By having views over shared data that include only lower level information, those applications that do not need highly classified data can operate at more economical levels. If data can be downgraded by interspersing a mediating security filter more operational efficiencies can accrue.
- **Database interface:** The relational model has provided a broad-based standard, although there are still many minor divergences among implementations. New approaches such as the Java Database Connectivity API are making it possible for non-relational systems to be "wrapped" to the relational model.
- **Information presentation:** Web technology has provided a commonality through ever-improved browsers that was impossible to foresee as recently as the mid-to-late 1990s.
- **Data representation:** While internal formats in computer software still differ markedly, web-based interchange formats such as XML and Resource Description Framework (RDF) are becoming widely accepted.
- **Data value semantics:** Developments leading towards a semantic web are promoting effective sharing of data element values, although we will likely see convergence on a domain by domain basis, rather than globally. Mediating techniques will still be needed for areas such as integrating information from military and commercial sources.
- **Application heterogeneity:** The many sites that share requirements and access to Air Force information systems also differ. The needs of a fighter wing overlap only partially with those of a

refueling fleet. While the common software modules should be universally available, there is little benefit in insisting that identical software configurations be used and maintained everywhere. Much waste can ensue from such forced commonality, since having to provide meaningless data will wear out personnel and reduce the accuracy and "up-to-date-ness" of relevant data.

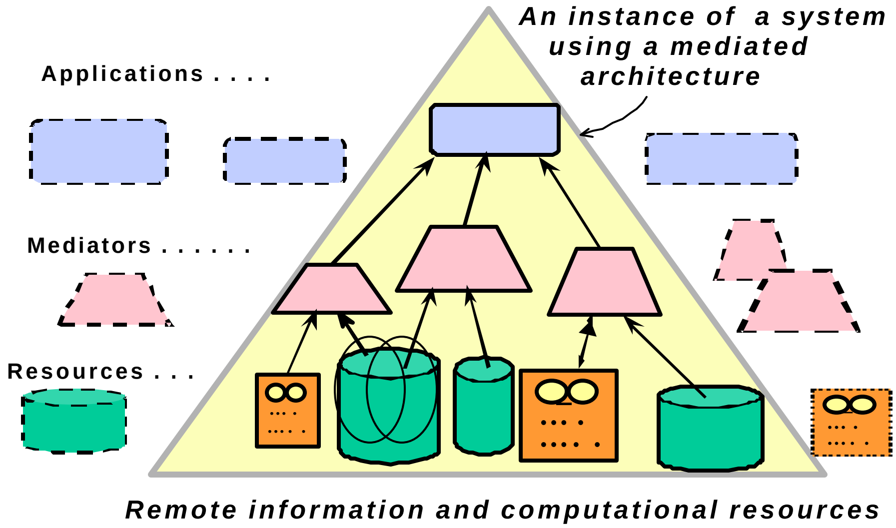

Figure 4-6. *Mediated Information System Architecture*.

Once heterogeneity of data sources is accepted as a way of life, and dealt with intelligently, many barriers to incremental progress in information system development disappear.

4.3.5 Migrating Applications

The contents of the individual databases are important sources of information when migrating information systems. Much data is currently replicated in Air Force systems. As the information in formerly stovepiped systems becomes a resource for all of the Air Force, much data replication (and the associated interface applications) becomes unnecessary. Replication by design still has a function for availability, risk reduction and backup. There are several COTS reverse-engineering tools which could extract information from databases and the source code of applications that use these data.

The most important aspect of a mediated approach is the fact that ongoing changes to system A do not affect the other systems at all, as shown in Figure 4-6. It makes no difference to the other systems whether their requests for shared data are satisfied directly from the shared database S, or through a mediator from A's database. This cuts the "you can't change anything until you can change everything" knot. Each system migrates at its own pace and a coordinated release and cutover are not required. Incremental changeover from obsolete systems to modern technology is enabled.

4.3.6 Modularity in Applications

To enable long-term effective insertion of improvements into Air Force systems, we recommend designing them with interfaces that allow incremental changes to the software. These principles are well understood in software engineering but are often ignored in DoD software acquisition where performance and cost are the only criteria, in spite of the fact that significant software efforts occur after delivery to the military.

4.3.6.1 Module updating and replacement

Both databases and applications accessing these data sources would benefit greatly if they had clear-cut programming interfaces. Most commercial software packages come equipped with an "Application Program

Interface" (API). For example, Oracle, Sybase, and virtually every database vendor supports various kinds of APIs. The API provides a set of functions that may be used by third party applications that use the package. The third party application communicates with another package through the latter's API. In short, an API is like (in a more sophisticated sense) an automatic teller machine (ATM). When we interact with an ATM machine, we do so in only a limited number of ways (e.g. withdraw cash, execute transfers between accounts, etc.). We do this without knowing the details of how these operations are implemented within the ATM. A third party application can interact in the same way with a package's API to access the functionality of that package without getting bogged down in the internal code. This mode of interaction offers huge advantages:

- First, the package can change and/or improve its internal components without worrying about adversely affecting applications that access it. To ensure this, the input and output types of its functions must remain unchanged, but the internal implementation of the function itself can be completely or partially revamped. This is a great boon because software programs are continuously upgraded, bugs are continuously fixed, and patches are provided on an ongoing basis.
- Second, if the package is being expanded to provide additional functionality, these new capabilities can simply be offered to applications as new API functions without affecting existing operations of either the package or the applications – in other words, the transfer can be significantly smoother than it currently is in most Air Force systems.

4.3.6.2 Moving modules to other platforms

A modular development of a server and/or application is also very helpful when the server/application is moved from one host platform to another. It is often wise to break down the design into clearly articulated components. All components fall into one of two categories – they are either "platform-dependent" or they are "platform-independent." In many cases, for instance, the *display* properties of a graphical user interface will be platform-dependent. As Graphical User Interfaces (GUI) typically allow users to express functions they want executed, the functions themselves can be implemented in a platform-independent way. What this implies is that if such a program is moved to a new platform (e.g. new operating systems [OS] or a new hardware platform), then the platform-independent components can be transitioned over with great ease (a relatively small amount of work may be required), while some re-engineering of the platform-dependent components may be required. It is therefore advisable to whittle away at components until the platform-dependent components are reduced to the smallest possible size – the larger and more complex they are, the harder it is to transport them to new modules.

4.4 Maintaining the Applications

Maintenance costs are often the dominant cost components of information systems. Software maintenance is needed when new hardware becomes available, when new information sources become available, when new technologies for fusing information are developed and when application users issue new demands. The cost of maintenance of any long-lived resource that must adapt to changing circumstances is substantial, be it airplanes or software. Computer hardware, fortunately, requires little maintenance today, because it is best replaced on a 3- to 5-year cycle.

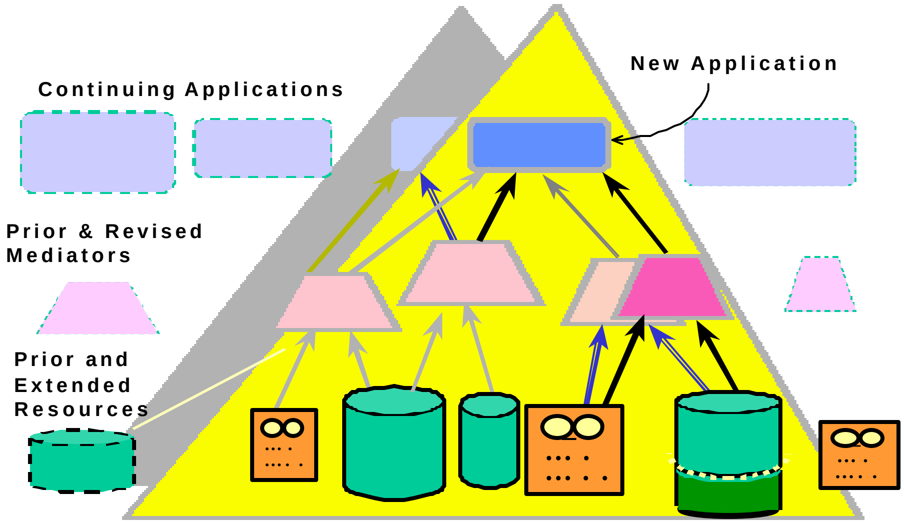

Figure 4-7. *Incremental System Maintenance*

Foregoing software maintenance, however, makes the information systems useless and even dangerous. Supplying wrong, obsolete, or confusing information to military decision makers creates high risks. Updating of information systems often also means that new data are needed. At other times new and useful data becomes available that motivates updating of applications. In either case, the incremental maintenance approach enabled by mediation allows rapid deployment of system enhancements while allowing continued operation of older applications. Some management effort is needed to ensure that obsolete applications are killed as soon as feasible. Maintenance costs are distributed over a longer period and systems can be updated selectively so that the quality of information systems can be kept high. To keep maintenance requirements visible in the acquisition process, it will be necessary to assign comparable values to maintainability and performance. While a poor approach to maintenance costs people and delays, pure performance can be improved by hardware, and is actually easier to deal with.

4.5 Interoperation

Incremental migration, as foreseen in this report, depends on the capability to link modules, subsystems and major systems rapidly and effectively without the loss of relevant information. We refer to the operation objective as “Interoperation.” Many modern technologies promise interoperability, but their applicability to the scale of Air Force requirements must still be ascertained. There are numerous issues that must be supported in order to achieve effective interoperation. These issues include the following items.

4.5.1 Wrapper Creation

Each data source must have a description of the data it contains as well as the assumptions associated with that data. This data (or “metadata”) often includes answers to the following questions:

4.5.1.1 What is the syntactic structure of the data?

This component is often referred to as the *schema* of the data. In many relational databases, the data will be stored as a table, where each row (or “tuple”) denotes a unit datum, and each column denotes attributes of the table. In such a case, the schema typically specifies the names of each column (e.g. LAST-NAME, FIRST-NAME, SALARY). With legacy applications becoming widely accessible via the Internet, however, many non-relational data sources have more complex structures. Such data sources may store information in a more hierarchical form using structures such as

```
<PERSON>
<NAME> <LAST-NAME> …</LAST-NAME>
            <FIRST-NAME> .. </FIRST-NAME>
</NAME>
<SALARY> … </SALARY>
<SSN> … </SSN>
</PERSON>
```

Such a description says that a class of items of interest called PERSONS is such that each member of this class has an attribute called a NAME, which in turn has two sub-attributes, FIRST-NAME and LAST-NAME. In addition, PERSONs have an attribute called SALARY (with no sub-attributes) and another attribute called SSN.

4.5.1.2 What is the semantic structure of the data?

The preceding question specifies the structure, but says nothing about the semantics. For instance, it does not tell us if the units used to specify SALARY are in US dollars, Canadian dollars or something else altogether. Wrappers must contain information describing semantic information about the syntactic schemas they describe. There are a number of commercial packages available for describing both the syntactic and semantic schemas of data.

4.5.2 Translation

When interoperating between a set of data sources, it is wise to adopt a uniform format within which the data obtained from these sources (either through importing or as answers to queries) can be stored. Once a mediator obtains data, that data must be translated into the internal format used by the mediator.

4.5.3 Carrying out Incremental Migration.

In today’s “24x7x365” world, it is impossible to shut down a system for any length of time. This means that when we add a new system, it must be seamlessly introduced into the interoperating federation with minimal glitches and loss of system usage. Incremental migration is well supported if the guidelines specified in this section are followed. Performing incremental migration requires some expansion to the mediator. First, it is possible to test the behavior of the expanded mediator (referred to as “New”) by replicating a copy of the old mediator and incorporating the desired updates to it. When the mediator is designed properly, then it is typically possible to have two copies of the mediator function simultaneously (old and new), with requests being processed by each according to a gradual ramp-up schedule (the percentage of requests processed by the NEW mediator is increased gradually till it equals 100%) until the old mediator is phased out. This strategy ensures that end users are completely protected from the transition.

4.5.4 Data Quality

Poor data quality can immediately cause a decrease in the confidence that users place in a system. Rule-based mechanisms can be used to automatically check the quality of data. Commercial tools for data quality that use rules to identify "dirty data" are well developed.

4.5.5 Security and Integrity

Two critical aspects of interoperability are security and integrity. When a set of data sources is distributed, each of the sources must be protected against attack. When the set of sources is integrated, however, the centralized integration platform also must be carefully secured. This is because an intruder who compromises the integration or mediation software potentially gains access as an "insider" to the data sources mediated against.

4.6 eXtensible Markup Language (XML)

XML is a family of new technologies and standards for web-based information management. XML technology makes many existing tasks – for example, creating, distributing, and processing an Air Tasking Order – much easier. Taking full advantage of XML requires more than technology insertion; it requires changes in USAF "business practices" and in the way we build information systems for C2 and combat support.

4.6.1 Role of XML

XML (eXtensible Markup Language) is a new standard through which data sources can describe their structure in a uniform way. This has great advantages. If all data sources use this structure to describe how their data is organized, then programs that access these sources can read these structures and can *automatically* read the data and split it into components. For example, consider a data source that contains employee information. It may describe its data as follows:

```
<PERSON>
<NAME> <LAST-NAME> …</LAST-NAME>
         <FIRST-NAME> .. </FIRST-NAME>
</NAME>
<UNIT> <UNIT-LOCATION> … </UNIT-LOCATION>
        <UNIT-COMMANDER>
           <LAST-NAME> Backus </LAST-NAME>
           <FIRST-NAME> John </FIRST-NAME>
        </UNIT-COMMANDER>
</UNIT>
<SALARY> … </SALARY>
<SSN> … </SSN>
</PERSON>
```

An application that reads such a description knows automatically that each object stored in such a data source has the above format. Hence, information about a single individual may be stored as:

```
<PERSON>
<NAME> <LAST-NAME> Doe </LAST-NAME>
         <FIRST-NAME> John </FIRST-NAME>
</NAME>
<UNIT> <UNIT-LOCATION> Uzbekistan </UNIT-LOCATION>
        <UNIT-COMMANDER> <LAST-NAME> Backus </LAST-NAME>
                         <FIRST-NAME> John </FIRST-NAME>
        </UNIT-COMMANDER>
</UNIT>
<SALARY> 56000 </SALARY>
<SSN> 201-88-7699 </SSN>
</PERSON>
```

As the application knows the format of the data, it can automatically read a stream of data, automatically identify its individual components and know how to reference the individual components for its own internal use.

4.6.2 Advantages of XML

The primary advantages of XML are:

- Applications can automatically read and process the data.

- Understanding of the data is often supported by the fact that communities in a given area of interest or market adopt standards for their particular community. For instance, processes are already being put in place to develop libraries of XML schema for use in the military (by DISA) and in the Air Force.
- Many C2I systems are already using XML, making it a leading contender for this kind of data description, and the commercial world is moving towards greater use of XML.

4.6.3 Disadvantages of XML

Though XML is clearly advantageous over current practices, there are also lingering questions about it. The first major concern is that legislation of content standards has never worked except in relatively small communities. While XML is thus advantageous in allowing common syntax to be used between many applications, it does not handle the semantic divergences between user communities. Thus, some in the Air Force believe that XML DTDs and Schemas are a "silver bullet" which solves all the hard problems in defining common data standards – but they aren't. The second major concern is that the robustness of the various programs that parse, process, query and otherwise manipulate XML data have not been fully tested, and there is much competition in the marketplace over tools with varying capabilities. Maturity of this market is coming, but it is still not yet here.

4.6.4 Using XML

People are able to examine and make sense of data given surprisingly few hints. We can look at the symbols on an airline ticket and decide when to leave for the airport, or at the columns of words and numbers on our bank accounts and understand our financial status. Computers do not have this cognitive ability. A computer system will not do the right thing unless the format and meaning of the input data are matched to the expectations of the software in meticulous detail.

The Extensible Markup Language (XML) is a web language standard for making documents "self-describing." XML is a set of rules for defining formats for data that are easy for programs to generate and process. These data formats can also capture some of the data's meaning. XML is therefore very useful for supplying data to people and machines across the web.

To better understand XML, compare it to its predecessor, the HyperText Markup Language (HTML). HTML is a language for *displaying* information on web browsers. The markup "tags" in HTML tell the browser software how the data should appear on the screen. This works very well for simple displays of airline schedules and bank statements. But advanced users want complex displays – of mathematical formula, chemical structures, musical scores – and no single, fixed language like HTML will satisfy the desires of all these user communities.

XML is a language for creating many customized markup languages. Users don't have to agree on one language; each community can have its own – while all still use the same browser software. XML accomplishes this feat by separating the meaning of data from the display of that data. The markup tags in XML are chosen to describe what the data is "about". A separate "style sheet" describes how the data should be presented on the screen. The same data can be presented differently, and tailored to different user needs, simply by writing another style sheet.

XML-based search engines can take advantage of the meaning captured in the markup tags to produce a much better, more focused search. Today, a web search for the key phrase "General Jumper" will return nearly 400,000 hits, because HTML does not allow us to distinguish between documents about Air Force officers, documents about parachuting, and documents that merely contain those two words somewhere in the text. Searches based on the meaningful tags in XML documents can avoid this problem.

XML data, with its well-defined format and meaningful markup tags, is very well suited for the exchange of data between computer systems. It is not necessary for independent systems to adopt internally the same data model and data element definitions, if they can both process the same XML data format. The use of XML for application data exchange is at the heart of many eBusiness approaches, and will eventually be more important than the original use of XML for flexible, powerful display of data to users.

4.6.5 Why Use XML?

XML is a very important development, but not because of *technical* innovation. The technical ideas in XML have been around in one form or another for many years. XML is powerful because it is a part of the "Internet phenomenon". XML is based on open standards. There are plenty of inexpensive (or free) tools. It is present everywhere the web is present – it will soon be a part of every web browser. There are many people with XML knowledge and skill, and more people learning XML every day. These factors combine to form a "virtuous circle", where each new use of XML reduces the cost and increases the value of the next. XML is powerful for *economic* reasons, not *technical* reasons. This makes a difference, because many problems that were technically hard before XML are still hard today.

4.6.6 What are the Technical Problems with Using XML?

Suppose you want to search for General Jumper's biography web page. Your search requests information about military officers, the web page provider marks the page as being about a military officer, and so your search returns one hit instead of a half-million. Problem: how did you know to use the term "military officer"? If you use different terms – if one of you chooses "Air Force personnel" instead – then the search fails. Advanced searching and related techniques like "publish and subscribe" depend on a *common vocabulary* of search terms shared by information producers and consumers. Building these vocabularies on the Internet scale is a huge challenge – and XML alone does not solve the problem.

For example, suppose you want your medical laboratory system to take orders from a hospital's information system and return the test results. This won't work unless you have a common set of definitions for things like "blood type" and "A-Neg". Establishing this sort of agreement is not too hard for any *two* systems. But what if your laboratory has many customers? What if the hospital uses many laboratories? Then you don't want to make agreements by pairs of systems; you want one agreed-upon standard for the whole community. Building these community data definitions is very hard work – and XML alone does not solve the problem.

For another example, suppose you are planning an air mission, so you query multiple information sources about friendly asset locations. The result tells you where the friendly forces *are*. But what you really want to know is where they are *not* – specifically, that there are none within the target zone. You can't get that answer unless you know whether your information sources, when added together, are *complete*. Reasoning about the contents of information sources in this fashion is difficult – and XML does not solve the problem.

All of these are *data management* problems. Solving these problems (and others) will require cooperation between the people who use and operate automated systems, the people who build the systems, and the people who decide what the systems are supposed to do. These groups of people must work together to control our data definitions, the contents of the databases that use those definitions, and the right to query and update those databases. XML does not solve these data management problems. If anything, XML makes their solution more important. These problems have always been the most challenging obstacle to seamless information. With the advent of XML, they are fast becoming the *last* obstacle.

4.6.7 What can XML do for the Air Force?

Many of the Air Force's existing "business practices" may be improved by inserting XML technology. For example, applying XML to the Air Tasking Order offers many benefits to an operational staff. Users would be able to search, extract, and display the information that is pertinent to their mission. We would have greater flexibility in the display of information – "What You See Is What You Want" instead of "What You See Is All You Get." The information could be made available in several ways: posted to a secure web site as an HTML file or passed to a telephone in audio format. All of these things can be done without XML. But XML technology makes these things *much easier* to implement – faster and cheaper, in one recent experiment, by a factor of 30. XML technology will also likely be at the heart of the Joint Battlesphere

Infosphere (JBI). The JBI will use XML as the basis of the “structured common representation” for collecting, organizing, and distributing information.

However, the enormous benefits of XML that are touted by the commercial world can come only through *changed business practices*, not merely by inserting XML into existing practices. XML technology only *enables* a change. The key factor is business practices that accept heterogeneity and encourage interoperability. Technology alone cannot solve these problems.

4.6.7.1 Use standard tools to convert database information to XML

The increasing importance of XML motivates many commercial suppliers to provide tools for making databases accessible via the web. While we do not yet see a single product that is fully adequate, there are enough tools so that an incremental conversion of Air Force data is fully feasible. However, no conversion process should be rigidly standardized now. Again, we will have to accept heterogeneity if the Air Force is to accept the benefit of ongoing commercial developments.

Metadata standards for XML are also still undergoing evolution. The original data tag definition (XML DTD) approach has been seen by computer scientists as too weak, having (for instance) no type capabilities. But, since it is fully adequate for documents, much XML data is still being described in this format, and many domain standards use DTD. The replacement, XMLschema, is now available and may supersede DTDs in due time, but we expect that both standards will be available for some time. In addition, the advent of web services and other new web technologies are increasing the need for better modeling capabilities, and languages like the Resource Description Framework (RDF and RDF Schema) are beginning to be used to augment XML-based systems. Again, tools for working with and processing these new languages are making great strides in the commercial world, and the Air Force should be able to take advantage of them to aid in data conversion efforts.

4.6.7.2 Wrapper technology

Obviously, it will be a long time before a large majority of Air Force C2 data resides in XML-based databases. Air Force information systems will have to access remote legacy databases of military, public and commercial origin. Here wrapper technology and mediation will have a role. Wrappers can be used to make information that was not meant to be shared or that was meant to be accessed only directly by humans available to remote computers. In particular, web information presented as friendly HTML pages should be automatically accessed by the integration of mediators.

These needs are not unique to the Air Force. Commercial services, such as comparison shopping sites, have developed many wrapping technologies. While some automated tools exist to help in the building of wrappers, there is typically also adaptation needed to deal with unusual formats, poorly defined semantics, etc. There is much commercial competency here, and a number of small commercial companies can respond rapidly to needs for wrappers. It will be important, however, to clearly specify the Air Force’s standards for acquiring such data. Both DTD specifications and XML schema specifications with good domain descriptions will be helpful here to allow rapid, incremental broadening of access to public and commercial resources. Emerging languages for the specification of vocabularies (RDFS) and ontologies (DAML+OIL) should also be on the radarscope of AF acquisition and development personnel.

4.6.8 What is the XML “Way Ahead” for the Air Force?

The most important first step is to collect and expose XML data definitions and schemas as they are created. Then people who want to understand an XML resource can obtain the documentation they need. Also, people who want to create new XML resources can discover and reuse existing definitions where possible. DISA’s XML Registry (http://diides.ncr.disa.mil/xmlreg) is taking the lead in this necessary first step.

We can expect the commercial world to continue developing XML technology and standards. The Air Force should simply follow these developments – there’s no need to be a leader here, and we shouldn’t try. Most of this work will be done by the commercial world, although there may be some special AF requirements that will not be met without AF participation in the process, either directly or through DISA representatives who are serving on some of the working groups.

The biggest part of the XML “way ahead” must be the transitioning of AF acquisition policies and AF contractors and system builders to the new “business processes” that XML technology will enable and their direct application to the C2 needs of the Air Force. The JBI is the "way ahead" to these new C2 business processes. The JBI is much like an “eBusiness” framework for C2 - it will allow us to quickly implement the new processes that XML makes possible. Success of the JBI depends on developing shared subject-area vocabularies and data definitions. An essential step is to revise data management policy and procedures so that proponents, builders, and users work together to produce the vocabularies that the JBI needs, and to track emerging trends in the business world so as to apply these to AF C2 needs.

An equally important step is to encourage (by funding) cross-program collaboration on these data efforts. Historically, we only funded individual programs and only got programs that worked individually. That approach is insufficient for eBusiness. Successful development of shared subject-area vocabularies and data definitions requires adequate funding of shared activities. This need for collaboration extends across the C2 and combat support environments and is most intense at the seams in the operational architectures where different functions interact. The recommendations section of this report describes some necessary changes to the organization of AF C2 acquisition efforts that include databases, and the implementation of AF C2 systems in the proposed structure will be made possible by XML and the other commercial technologies relating to it.

4.7 Warehousing and data analysis

Warehousing is a technology that aggregates historical data from many sources and moves it to a single large data server. They are well developed in commercial practice and serve large-scale On-Line data Analysis Programs (OLAP).

4.7.1 Warehouse Implementation

The data in a data warehouse is conceptually represented as a multi-dimensional matrix, with axes such as time and date, location, customer category, product types and the like. The actual labeling of the axes is an implementation decision, so that parameters useful for Air Force applications can also be inserted. Data warehouses are updated by the acquisition of data from many sources. Where source formats differ, techniques that support mediation can be employed. The data stored are often abstracted or summarized to provide a granularity that is adequate for subsequent analyses, and sufficiently reduced for effective processing and long-term storage.

Keeping warehoused data current is a major issue. The two strategies that are most used are periodic rewriting of the entire warehouse content from the sources or periodic updating of content from recent information. Which strategy and update frequency is chosen depends on several factors:

- The length of the historical view into the past needed by the warehouse users.
- The ability of the sources to keep historical data adequate for the warehouse.
- The degree of transformation for the selected granularity.
- The cost and load on the sources for data acquisition.
- The costs and benefits of keeping the warehouse and its users current.

Because of the massiveness of data warehouses, they are typically not kept up-to-date in real time. Typical update frequencies range from the daily to the quarterly, depending on the volume and benefits to be gained from current information.

Applications where warehousing solutions are appropriate to Air Force Information systems are typically centered around intelligence applications. Some limited applications are to support locations that may become isolated, so that the warehouse serves as backup for such instances and for data that arrive at different moments but must be viewed consistently, such as images of moving targets that are obtained by different vehicles and modalities. Any legacy application that falls into the range of capabilities provided by today's data warehousing technologies should be moved to commercial technology. Since the commercial products are improving, there must then also be a budget to acquire those improvements as they become available.

4.7.2 Analyzing Data from Warehouses

Current data analysis tools range from browsers that allow rapid viewing of the warehoused data in aggregations (over the defined axes) to powerful statistical tools. Some novel techniques (usually referred to as "data mining") support discovery that is not based on accepted statistical models, where independent and dependent variables can be predefined. A weakness of the avoidance of statistical or causal models in OLAP is that secondary effects are hard to discern, since they will be lost in the noise associated with primary relationships. Those primary relationships are often already known to the analysts, so that the discoveries made, even though automated, do not convey novelty. We do expect the commercial world to continue to invest in these technologies, in part to justify the investments already made in warehousing technology. That means that innovative OLAP technology is bound to improve. Where military analysts already have good first-order models, the available statistics-based tools will serve them well.

4.7.3 Warehouse Interfaces

Where warehouses are closely coupled to analysis systems, it is best to follow the existing technologies. SQL is often used, although sometimes with extensions that go beyond the published standards. SQL is of course an available interface standard whenever commercial databases are used for high-demand information sources such as data warehousing. The large amounts of storage that warehouses employ make use of XML as an internal representation method infeasible. If, however, warehouse technology is to be integrated into Air Force Information systems, they should export XML representations. Again, this will require that Air Force data descriptions be made available in an XML metadata standard, as DTDs or XML schema.

4.7.4 Warehouse Summary

Data warehousing and the associated analysis programs are rapidly maturing technologies. The technology is relatively heavyweight and will be useful for the Air Force in some situations, but certainly not all. The delays implicit in obtaining data from many sources, summarizing it according to the n-dimensional matrix chosen, and inserting it while maintaining global and temporal consistency make warehousing inappropriate for supporting operations at near real-time speeds.

4.8 Data Security

Security of critical data is a major issue in Air Force information systems and is handled more consistently than in the commercial world. Traditional security protection is predicated on the existence of document classifications and ensuring that accessors are authorized to access documents having that classification. The model for secure processing has become increasingly sophisticated. Distinctions are now made of several mandatory and many categories of discretionary classifications. Categories include types and categories of potential accessors, such as "No Foreign." Rules are defined to establish safe mandatory levels for information based on multiple sources by setting the results to the most restrictive level. Unauthorized access is prevented by data encryption, both in storage and during transmission. While multiple or distinct encryptions are feasible for distinct classifications, these are not used in practice because of the complexity they induce, so that once access is authorized, the classification label must be checked to assure protection.

When adapting commercial software, the security capabilities in its intended setting must be reviewed. Since security regulations deal with the protection of data from unreliable input and the direction of output to only authorized personnel, the validation concentrates on the interfaces. Software that is obtained to operate at a single level of classification and runs on subsystems that have no inputs at lower levels and no outputs at higher levels can be installed with minimal precautions.

Each subsystem must protect itself, which means that it is, directly or indirectly responsible for:

- *Authentication* – validating the user identification.
- *Authorization* – assigning rights to the identified user, typically by role.
- *Labeling*—making sure that all outgoing data is identified with the classification level of the subsystem.

If data of lower classification level are introduced in such a system, they automatically become system-high. Data at high classification levels are more costly and are often processed slowly, since fewer personnel resources will be available. It is thus desirable to provide for multi-level data management wherever possible

4.8.1 Distributed versus Single-System Multi-level Security

In installations that operate on multiple levels, much care and validation are needed. Multi-level secure operating systems and multi-level secure databases are available, but falling out of favor. Their cost is high for two reasons. Because there is no demand for them in the commercial sector, there is an additional cost to make them secure and to validate their operation. The time delay is the reason for the second cost factor: these systems are usually a year or more behind the technology curve and are consequently based on systems that were costlier or performed less capably.

The low cost of modern hardware leads to configurations where subsystems are partitioned to operate in a single level of security classification. Now the protections required for multi-level security requirements are placed on network interfaces. Data classified at a low level can be shipped redundantly to subsystems at its level and to subsystems at higher levels, if needed for integration at those subsystems.

4.8.2 Problems due to Collaboration

In the post-Cold War world, the rapid assembly and termination of collaborations, often with former foes, presents a new set of problems. The traditional requirement for secure operation is that all documents have classification labels. To predict, at the time of document creation, all possible future uses of the information was difficult and is increasingly impossible. We engage in instant collaborations and develop many specialized technologies, so that the web of possible classifications becomes dynamic and unpredictable. Having a complex classification scheme also leads to an increasing risk of errors, which is counteracted by choosing restrictive levels of classification. Now information is likely to be withheld from users who should be able to utilize it, leading to inefficiencies and operational risks.

4.8.3 Down-level Information Transmittal

Data coming from higher levels cannot be shipped to systems at lower levels unless filtered. Furthermore, data can never be shipped among discretionary levels unless already labeled with the destination tags. If that tag is a new category, any automatic transmission of such documents is disabled and the filtering of the contents is required. Today, such filters are mostly mediated by human personnel who have access to subsystems at more than one classification level. This creates a secure “air gap.” There is some automated filtering technology available, but since it must by its very nature be paranoid, such filtering software must be carefully validated.

A solution is to provide mediating nodes in the network specifically for the release of previously classified data to users who have legitimate access that is not properly reflected in the classification labels. Allowing such release requires actual checking of the data, i.e., its content and not merely its labels. Automation of such checking means that the entire message of the text must be checked for the exclusive presence of terms that are valid in the permitted context. The dictionary for such checking can be created by processing a number of documents that were verified by human inspection to be acceptable for release. Even an automated system must be augmented with human backup, since false indications of failures are likely.

Semi-automated checking, which is immediately feasible, requires that a security officer be notified whenever a document being checked for release contains terms not in the dictionary, and then augmenting the dictionary of the reviewing security officer. In a short time, the dictionary will grow so that routine documents will not engender a request to the security officer. While checking of contents consumes many more computing cycles than label checking, the capacity of modern computers makes it feasible. Modern spelling and grammar correctors work on the same principle.

4.8.4 Data Security Summary

The provision of security has a substantial cost. Exploitation of the relatively low cost of hardware can provide the isolation and redundancy that lower the cost of providing secure systems. Despite this, security is never absolute and tradeoffs may be necessary. If the provision of security to gain an extremely low level of risk inhibits the use of information and the use of software that will protect our soldiers and facilities from unexpected attacks, then the balance has swung too far. It will require much wisdom of the sort not easily encapsulated in formal directives to strike the right balance. New technologies for tagging, sorting, filtering and integrating data may be a help in making it possible to do an improved job of providing security without sacrificing flexibility.

4.9 Agents

There is now a growing body of work on software agents. Several impressive platforms for the creation and deployment of agents have been developed. Agents built using these platforms provide a variety of services, including the identification of interesting newspaper articles, software robots that perform tasks (and plan) for users, content-based routers, agent-based telecommunication applications, and solutions to logistics problems. More recently, agents have been used not only to integrate multiple legacy databases, but also to tie them to third party legacy and/or commercial programs.

A number of different classifications have been used to attempt to define an agent. Generally, an agent is said to be a program that satisfies three general requirements:

- *Adaptivity*: Agents must be able to adapt to changes in their environment.
- *Autonomy*: Agents must be able to take actions automatically without requiring constant human direction.
- *Collaboration*: Agents must be able to collaborate with other similar agents so that a federation of such agents can achieve more than the sum of their parts.

Techniques exist now to take legacy software and "agentize" them. In such efforts, one starts out with the legacy program's application program interface (API) and adds new structures to it. These new structures include (amongst others):

- A set of actions for interacting with other.
- A messaging capability.
- A set of rules that specify what actions to take in response to a change in the agent's environment.

For example, if one wanted to integrate two AF databases, one could add such structures to each of these two databases. The first database may be augmented with actions such as "Execute-SQL-Query", "Modify-SQL-Query" or "Send-message." The second AF database may also be augmented via such actions. The rules that these database agents use may vary. The first may have a rule saying that all requests about a project from a user (or an agent working on his behalf) should be answered as long as the user is a team member for that project. The second database agent may use a different rule – for instance, providing a lower level of service to the user's agent (or alternatively, the second agent may send a separate authentication agent a message with information about the request and may answer the request only after it receives the authentication agent's consent). Once the agent has created answers to its query, it could send its answers to a PowerPoint agent that creates a PPT presentation for the user, who visualizes the answer to his query using PowerPoint.

Agent technology has several advantages for integration:

- As each source/program is "agentized," it can participate in many "federations" which require its services.
- Changing the behavior of an agent (how it makes decisions in a given situation) is easy, as these behaviors are generally encapsulated in a special part of the program (an agent wrapper) separated from the base functionality.
- Agents can take actions such as replicating themselves dynamically when performance degrades to enhance performance (scalability) as well as reliability.
- Agents can be easily verified (as the number of rules in them tends to be fairly small even in cases where a monstrous amount of data is involved) – in short, agents embody the very principle of modularity by providing a few, clearly verifiable services.
- Agents are easy to maintain.
- Agents can easily be replaced by other agents in a federation.

However, agent technology also suffers from disadvantages:

- Agent technology is not as mature as other integration methods such as mediation.
- Though several commercial tools (IBM Aglets, Oracle Mobile Agents, Mitsubishi Concordia, ObjectSpace Voyager) exist for agents, most of these tools focus on providing mobility services and not services related to building methods to facilitate collaborations between legacy codes.

There are ongoing experiments on the use of agent technology in the DoD, and we recommend that the outcome of these experiments be closely monitored by the Air Force so that a positive outcome in the experiments can be readily leveraged.

A limited form of agent capabilities, that of fusing existing information based on a set of contextual rules, was identified as a key component for the development of Air Force C2 systems in the JBI. These "fuselets," as they were named, were not intended to be general agent programs, but smaller and more focused capabilities that could be easily and widely deployed. Fuselets were particularly intended for the integration and rollup of data from different legacy data sources and thus can serve to do the sorts of data integration tasks described in this C2 Database Interoperability report. Fuselets are an active area of AFRL/IF JBI exploration, and are expected to be deployable in Air Force systems in the near future.

4.10 Conclusion

Much progress will be made in the area of DB migration by simply putting some new processes in place, and this report makes a strong case for that. While many technologies are available in the commercial world, some of the needed processes will be either impractical or impossible to implement without continued development of technologies. In addition, there is a clear need to continue the development of technologies that will enable new architectures for dealing with data and information.

We have presented a technical migration strategy that moves Air Force information systems to a general architecture that enables sharing, use of commercial standards and tools, and incremental maintenance. Furthermore, we recommend a process in which systems and their applications are gradually converted to clients of local and remote shared data servers by using reverse engineering tools to develop the shared data schema.

The salient points of the resulting systems include:

- Use of XML and related technologies for the representation of shared information whenever possible.
- Let user interfaces be browser-based wherever practical.
- Employ commercial tools for data mining and performing other analyses.
- Insertion of data mediators between existing applications and shared databases.
- Accessing remote legacy or modernized resources through information integrating mediators, with an eventual migration to agent-based techniques as they become commercially available.

These technologies will allow the incremental independent updating of legacy application code to use the sharable resources and the incremental insertion of modular applications as needs or new intrinsic capabilities arise. During acquisition, aspects of operability and maintainability must be considered:

- Security provisions that provide adequate protection without hindering daily operations.
- Adherence to commercially accepted standards.
- Modular software composition so that systems can be flexibly configured and incrementally improved.
- Acceptance of heterogeneity in software and hardware at the system and application levels.

Despite the enormous gains that can be achieved by the immediate use of commercial methodologies and tools, the Air Force must still remain invested in S&T related to C2 systems integration and interoperation. Research and development specific to Air Force needs are required to help evaluate the tradeoffs made when moving to COTS-based software. Our recommendations in that respect are based on observations and experience, not on validated models. Integration technology has concentrated on commercial needs and thus AF-specific needs have not been explicitly addressed. For example, the extent of redundancy needed in military situations, where the loss of entire nodes and inter-node communication links can occur, will require performance characteristics that are not widely explored in the commercial world. The scalability of some of the XML technologies is also uncertain, as large-scale metadata management has not yet been proven. Security research has to consider the complex cases of collaboration, and not focus on making a simple good-guy/bad-guy model absolutely secure.

Chapter 5

Models For Successful Air Force Migration

Slide 35: Outline of Industry Practices

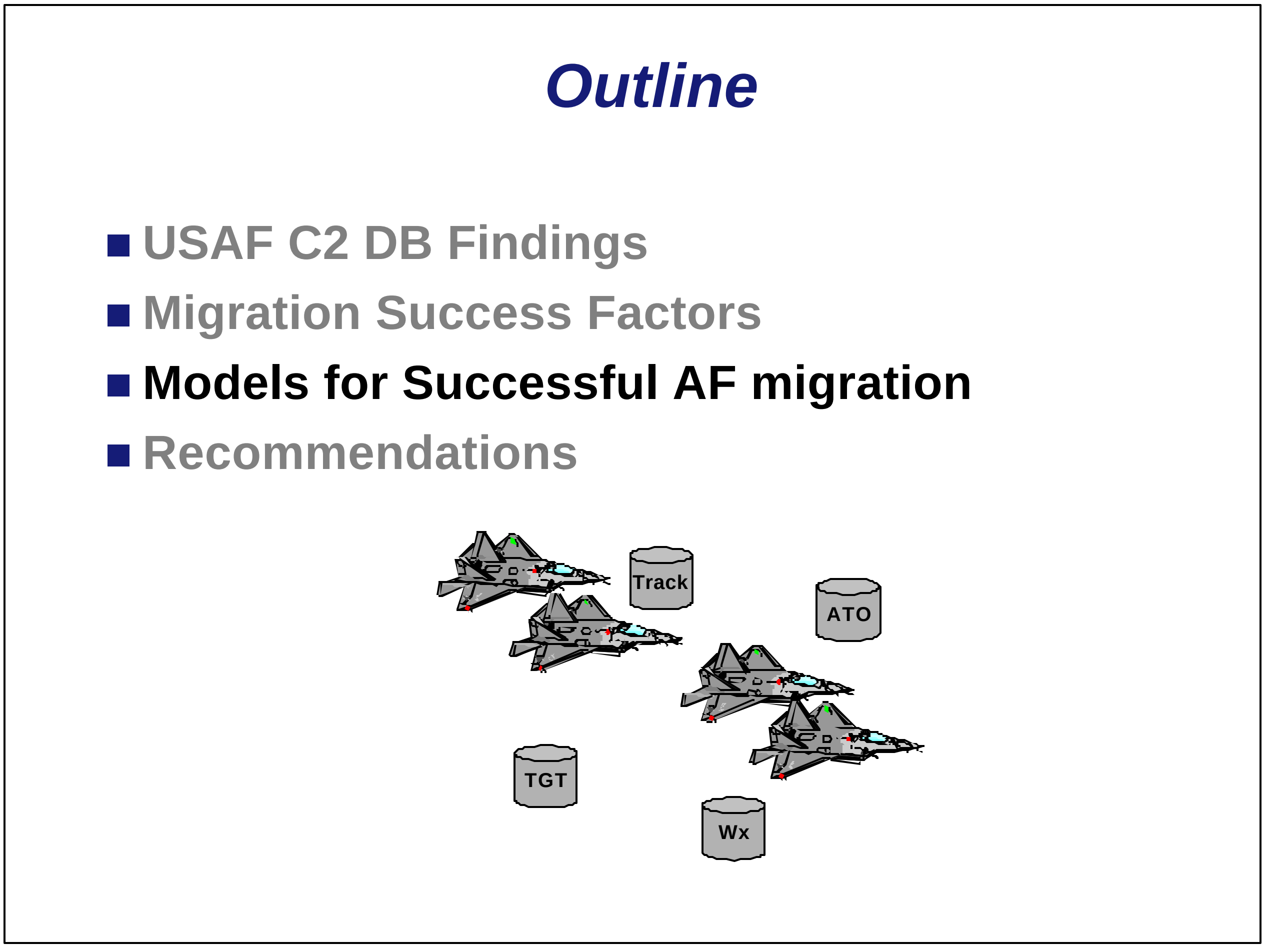

In this section, we discuss some of the ways that industrial best practices might be put into effect in the USAF.

Slide 36: Getting There From Here

Getting There From Here

- **The USAF can adapt the model used by corporate America to develop a C2 data management process**
  - **Assure the continued viability of the data contained in the legacy databases**
  - **Migrate the databases to a Joint Battlespace Infosphere environment**

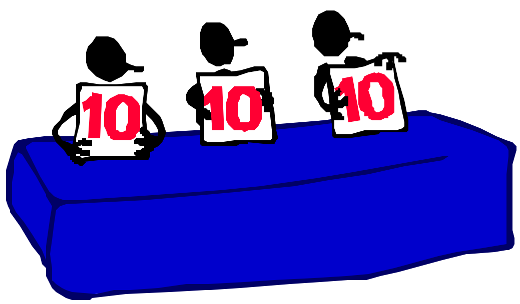

In this chapter, we discuss the manner in which corporate practices may be adapted to Air Force needs. The major lessons we draw from the commercial sector are the following:

- The Air Force must treat command and control data as a strategic asset and manage it at the enterprise level. This is to be contrasted with current Air Force programs that treat data as an asset of a particular system or command.
- The Air Force must adopt an evolutionary migration strategy for dealing with command and control systems, the databases that they use, and the CONOPS that define the operational use of these systems. This is to be contrasted with the current process of acquiring information systems, which resembles that used to acquire major hardware items such as weapons that involve a large development effort followed by a sustainment phase.

This section describes some of the special challenges the AF faces and the ways in which they may be overcome. The next section of the report turns these into specific recommendations. The challenges which the Air Force faces in migrating command and control databases go well beyond the merely technical. Three challenges stand out:

- Over the years, the Air Force has institutionalized a set of processes for system acquisition. This has resulted in stovepiped systems. The data is viewed as an asset of the system, not of the enterprise. This culture is ingrained in the acquisition process, the budgeting process and the various using commands. Any evolution toward a migration strategy will require a corresponding evolution in the culture of acquisitions.

- There is no enterprise-level focus on the information systems that support command and control. The idea that we will treat data and information as strategic assets means a corresponding requirement for top-level oversight of the migration process. We refer to this role as the "Guiding Hand."
- A particular challenge for the Air Force will be to recognize database migration as a continuous evolutionary process that must be supported by a budget permitting continuous system enhancements. We currently have alternate cycles of system acquisition (R&D funds) and system maintenance (O&M funds). The distinction in the "color of money" is a major barrier to continuous system evolution.

The main focus of this section will be a description of the specific actions the Air Force should take to facilitate the migration of C2 databases. These actions are briefly summarized below:

- Data must be recognized and managed as a <u>strategic asset</u>. Key points include:
  - Data is a <u>shared</u> resource.
  - Someone must be <u>responsible</u> for all data, and it must be possible to trace the pedigree of all data.
  - Data exists at <u>many levels</u> and in many forms in an enterprise (e.g. in legacy systems, databases, metadata descriptions, images, sensor systems, etc.).
  - Data must be <u>protected</u>.
  - Data must have a <u>confidence</u> factor. Data ages, and it may come from questionable sources.
  - Collection systems have inherent errors.
- Databases and their supporting systems must evolve. In fact, it is essential that the operational architecture co-evolve with the system architecture. They must co-evolve because as missions and weapon systems change, the supporting Command and Control systems must evolve to support the new CONOPS. This includes ensuring that the right data is available to the warfighters. The link between the operators and system developers is the Information Exchange Request (IER).
- Exploit prior recommendations. The C2 Database Migration study reaffirms and extends past Scientific Advisory Board recommendations. The JBI studies of 1998 and 1999 provided a framework for managing information to ensure that the right information reaches the right person at the right time. Sharing and interoperability are key features of the JBI. The 2000 SAB study on "Air Force Command and Control: The Path Ahead" emphasized the need for an Air Force focal point for command and control. It also emphasized the need to evolve systems to meet operational requirements and to take advantage of technology advances. The recommendations from this C2 Database Migration study are aligned with each of these previous studies.
- Equip the C2 data enterprise. If our key goals of treating data as a strategic enterprise and co-evolution of systems, databases and CONOPS are to be met, there must be major changes in the acquisition system, in the budgeting process and in the overall management of data assets. Many of these recommendations have been made before, but the Air Force continues to face a number of the same problems. Cost is often cited as a barrier to some of the recommended changes. Experiences in the commercial sector and at USTC show that by appropriately freezing some functions and reallocating the resources, it is possible to evolve data systems without significantly increasing costs. Sharing of data will reduce enterprise costs, but it may put a greater burden on some organizations. Moving away from the "big bang" model of system replacement and substituting the incremental evolution of components will result in more uniform spreading of costs over time. It will also provide fewer disruptions in the operational systems. All of these will require a guidance and authority that extend across the C2 enterprise. This is the concept of a "guiding hand" that will be expanded in the next section, where we make specific recommendations for the USAF.
- Train for the C2 data enterprise. Today our warfighters need to understand the CONOPS and the systems (both weapon systems and their supporting command and control systems) they use to execute an operation. They receive extensive training to prepare them for their missions. In the future, the training for C2 operators must also include an understanding of the data and the implications of its use. This is meant not to replace the people who maintain the database and

communication systems, but to understand the content and the implications of data use. As we evolve toward the JBI, there will be greater sharing of data and a greater dependence on the use of automated tools to assist in decision making. Knowledge of the data and its limitations will be essential to future missions.

Slide 37: The Migration Path

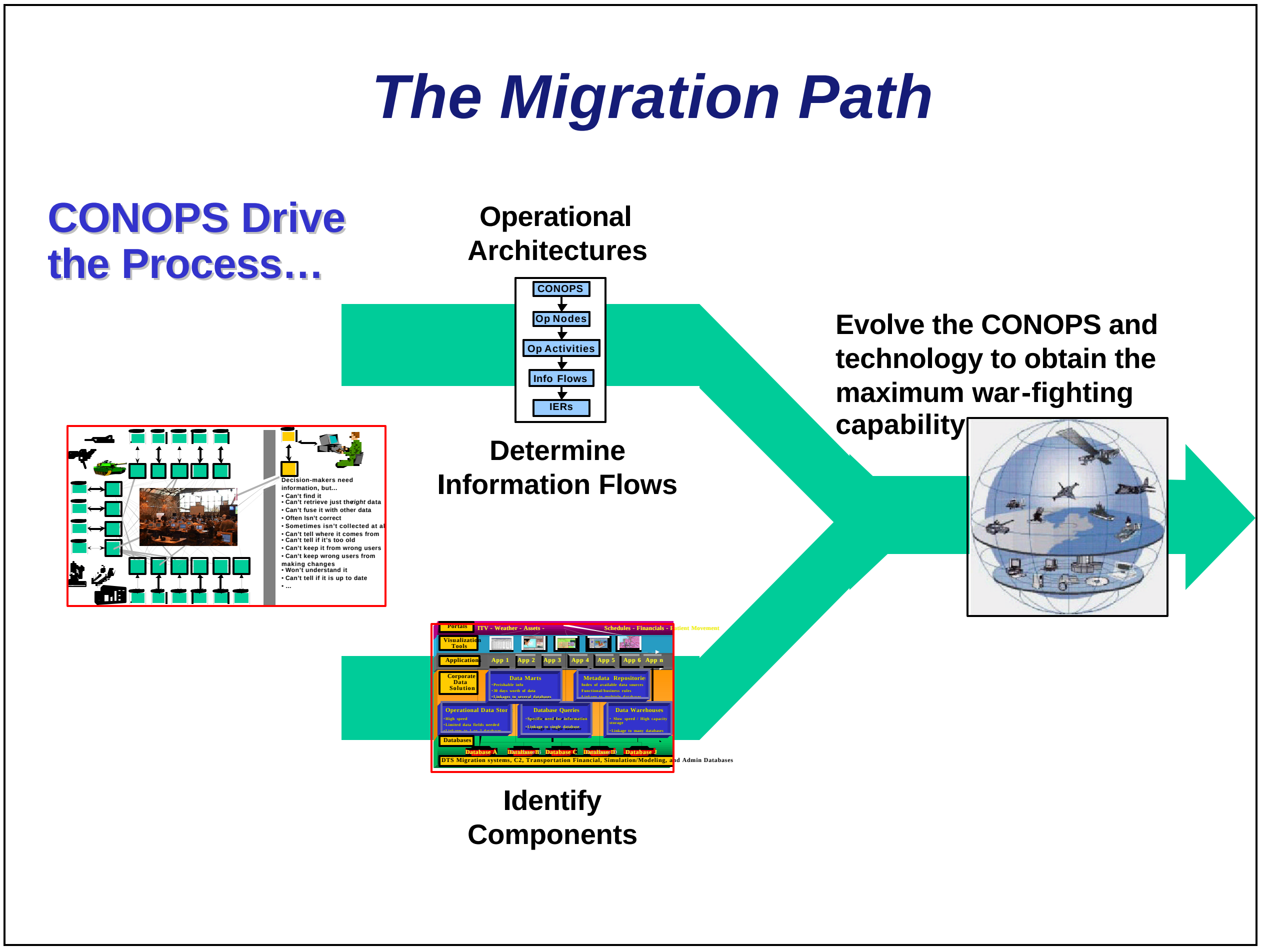

As established by the review of corporate databases and information management systems for decision-making, databases (and the activities they support) must evolve in order to be successful over the long term. This finding leads to two key questions: the first being, *what drives the evolutionary process?* The second is: *what are the existing technologies to bootstrap the evolution?* This study concludes that CONOPS should be the driver for evolution, with tools such as Information Exchange Requirements (IER); these are already used for data integration and they provide implementation leverage.

The role of CONOPS as a driver for C2 Database Migration is illustrated above. Migration consists of two concurrent activities: determining information flows and identifying the needed components. As information flows are identified, the needed components (either tools or structures) for effective information management become apparent. As new components become available from business applications or through DoD-specific funding, the information flow can occur in new ways. Together, the continued interplay between "what I need" and "what I can do" will maximize warfighting capability.

Slide 38: C2 Operational Architecture

C2 Operational Architecture

- **The operational architecture view is a description of the tasks and activities, operational nodes and elements, and information flows required to accomplish or support a military operation**
- **An operational architecture is derived from doctrine and/or an operational concept**
- **Operational nodes/elements are defined; operational activities are assigned to them**
- **The operational activities require data from other operational nodes/elements**
- **These information flows between operational nodes induce Information Exchange Requirements (IER)**

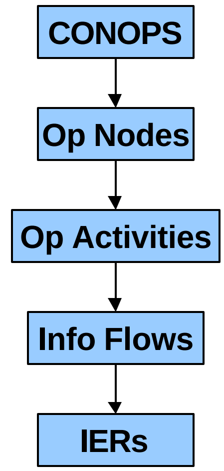

… to Produce Information Exchange Requirements

Mechanisms for determining information flows exist and can serve to bootstrap the evolution process. Perhaps the best-known mechanism is the construction of an *operational architecture*. An operational architecture represents a system in terms of its operational elements, activities, tasks and exchange requirements. As such, it is different from systems architectures, which describe the physical connections between subsystems, or technical architectures, which provide implementation details.

An operational architecture is derived from doctrine and/or an *operational concept*. Once the operational concept is extracted, *operational nodes/elements* are defined; *operational* activities are assigned to them. The operational activities require data from other operational nodes or elements. These *information flows* between operational nodes can be codified as IER. The use of IERs in resolving the gap between operator responsibilities and integrator responsibilities has already been noted in the SAB's 2000 study on "Air Force Command and Control: The Path Ahead." The bridging of this gap is critical to the evolutionary cycle, where operators focus on CONOPS and the overall systems architecture, which leads to a set of desired capabilities that are supplied and supported by the integrators. These capabilities are manifested as system updates, which then complete the cycle by leading to new capability needs that will be identified by the operators.

This study recommends that the cycle of evolution begin by pushing the rapid development of a USAF-wide C2 operational architecture, rather than restricting the focus to a single, albeit important, application. By deriving C2 data CONOPS, the immediate goals of a migration process are established and understood, thereby activating the evolutionary pressure to find or develop tools and processes to achieve these goals.

Furthermore, by examining migration from a USAF-wide perspective, spiral development of JBI services and the framework for the C2 operational architecture and experimentation would be enabled.

Slide 39: IERs in Evolutionary Data Integration

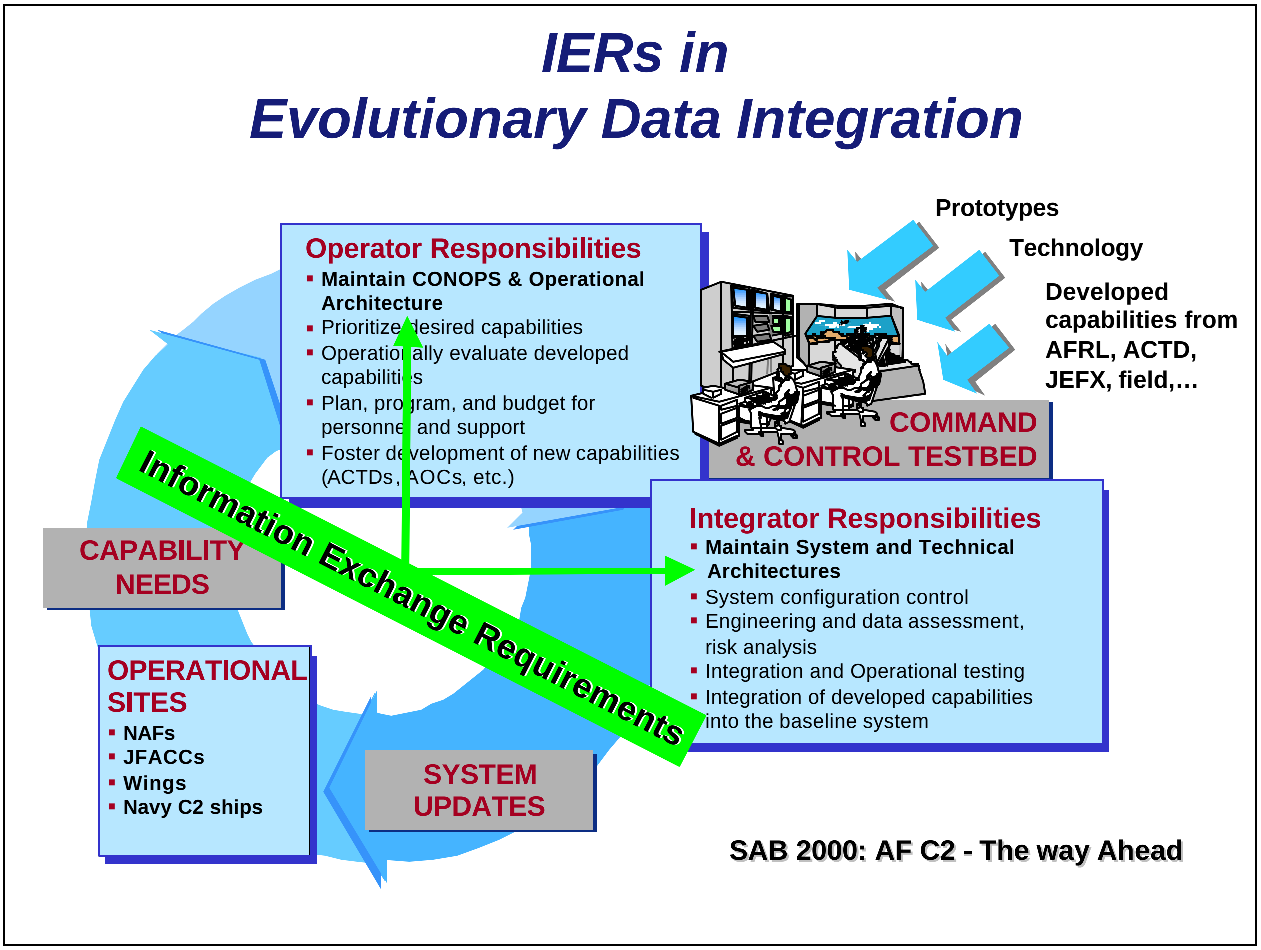

The 2000 SAB summer study was redirected to command and control because of the continuing challenges the Air Force faced in providing timely and effective command and control of its operational forces. There have been a number of similar studies over the past 15 years, and many of the 2000 SAB recommendations have been made in prior studies. The recommendations of the previous SAB study are very relevant to our study, and in fact, the two studies compliment each other. The essence of the 2000 SAB recommendation is to provide a focus and follow-through on C2 issues from a very high level. We summarize the findings of that report and the manner in which they relate to our study. We particularly emphasize the first bullet below, which is the focus of this slide.

- *Institutionalize an evolutionary C2 integration process.* Another recommendation of the 2000 SAB study was to create partnerships between operators, developers, integrators and operational testers. This was to be achieved by the creation of a spiral process aimed specifically at improving AF C2 systems and processes. We believe this spiral is also the place where we can co-evolve CONOPS and data sharing- that is, this evolutionary loop is a place where the IERs that come out of the process described in the previous section can be refined and tested.
- *Establish a single C2 manager at the Air Force level.* This is consistent with the C2 Database recommendation that there be a "guiding hand" at the CIO level of the Air Force to provide an enterprise-level focus on data and to treat data as a strategic asset (We note that in December, 2001, the Air Force was reorganized to create a deputy chief of staff for warfighting integration [AF/XI],

who will have enterprise C2 oversight). We believe this is an important step toward the implementation of our study results.

- *Field and evolve the Theater Battle Management Core System (TBMCS).* "It is time to accept the system and to accept the fact that continual upgrades will be needed to meet operational requirements and technology advances; the upgrades should be so planned." Within the TBMCS system are numerous databases that must be migrated over time. Many of these databases are maintained by other organizations (such as the intelligence databases, personnel databases, logistics databases, etc.), and it is important that they be maintained and evolved to meet the operational needs of the Air Force. A long discussion of TBMCS was presented in the "Current C2 Database Report" section provided earlier in this study.
- *Enable and encourage rapid technology insertion.* In particular, CONOPS and desired operational capabilities should drive the process. It has often been difficult to integrate new technologies rapidly into operational systems. The delays experienced in the integration of new technologies into DII COE provide a good example. To alleviate this problem, the 2000 study recommended the establishment of a C2 test bed to foster rapid development and integration of new technologies. Should such a test bed be established, it could also perform a valuable service in demonstrating processes for migration of databases.
- *Staffing and training must be consistent with the importance of C2.* Derive training requirements from CONOPS and elevate the stature of and advancement opportunities for C2 warriors. As is pointed out below, training for C2 data operators is an important recommendation of the C2 database study. Not only must data operators know how to access and use the various databases, they must also know the limitations of the data and the databases. The C2 operators must also be trained to share data and make their information available to others who need to access it. We will return to this issue later in this report.
- *Achieve information interoperability for the warfighter through the JBI.* This study urged the Air Force to seize the initiative to evolve the JBI as the basis for true interoperability. We concur, and our next slide explains why the JBI is the right environment within which many of the C2 database migration recommendations must be implemented.

Slide 40: Migration to the JBI Environment

Migration to JBI Environment

- Develop USAF-wide C2 operational architecture
  - Derive C2 data CONOPS
  - Produce migration path from legacy systems (e.g. TBMCS)
- Co-evolve JBI services and C2 operational architecture
  - Spiral development and experimentation

JBI

Promotes sharing (e.g. publish, subscribe)

Provides services for evolving a component-based architecture

Need a program to migrate AF C2 systems

A critical recommendation of our report is that the time has come to fully fund the JBI project, particularly the investigations into the development of JBI services at AFRL/IF. Below we summarize the reasons why the JBI is so important to AF C2 DB migration processes.

The Joint Battlespace Infosphere can be viewed from several different perspectives:

- It is first of all a vision of information management in support of the warfighter in which the right information is provided to the right users at the right place and time.
- It is an infrastructure that presents a set of information services to users and applications that are robust, policy-compliant, easy to use, interoperable across communities of users and transparently evolvable with the progress of technology implementation.
- It is an architecture that integrates innovative technologies such as "publish," "subscribe," "fuselets," "collaboration technologies" and advanced user interaction technology to continuously exploit evolving technologies to meet changing mission requirements. The JBI should leverage technologies from the commercial sector.
- It is a weapon system—actually a system-of-systems—that collects, integrates, aggregates and distributes information to users at all echelons. As a weapon system, it must be procured, evolved and managed, and this implies a process and associated budget to keep the system current, much as an aircraft system must plan and fund for block improvements. In the case of software, the specifics are less predictable in advance, but no less predictable in the aggregate.

In all of these perspectives, the JBI is addressing the issue of information management in a joint battlespace context and relies consequently on data from a variety of sources. While our recommendations on database migration must first be implemented within the current command and control environment, we believe that the JBI framework will greatly facilitate the implementation of our recommendations and that our recommendations will further strengthen and aid the evolution of the JBI. That is, the JBI is not a system to be acquired, but rather can be defined as a major step on our data migration path - the state toward which we wish to migrate many of our systems.

The JBI is the right environment for our recommendations:

- The JBI focus is on information management, with particular emphasis on sharing and interoperability. Since data sharing is a key goal of the C2 database study, technologies developed for the JBI will facilitate the sharing we feel is necessary for database migration.
- The JBI is a component-based architecture that exploits commercial technologies, including web-based interactions, standards such as XML, and commercially available components. There is a separation between components and the data they use. This architecture leads to a focus on the services to be provided rather than on a single system that integrates functionality and data.
- For future C2 DB migration, we recommend that the Air Force maintain an organic capability to integrate, test and evaluate the emerging data-related technologies coming out of the commercial sector. The JBI activity at AFRL/IF is clearly the leading activity within the USAF and it is important that this activity be aligned with the results of this study.

The database migration recommendations complement the JBI:

- The database migration study stresses the need to treat data as an enterprise asset that goes across both systems and organizations. This study provides specific recommendations on the establishment of enterprise-wide policies and responsibilities for data management. This study thus provides additional insight into a major area the JBI must accommodate.
- This study stresses the need to develop the CONOPS concurrently with the systems that support command and control operations. This includes making sure the necessary data is available in a timely manner.

A major theme in the JBI recommendations was that we should migrate current systems to the JBI. We also recommend such a migration strategy. This strategy is necessary because of the need to have an operational system at all times and because the cost of developing a replacement system can be prohibitive.

Slide 41: Recommendations

Manage C2 Data as an Enterprise Asset

- **C2 data must be treated as a USAF strategic asset**
  - **Warfighters' lives depend on decision-quality information**
- **Centralize corporate data responsibility**
  - **Provide enterprise-wide tools and policies**
  - **Support and consult with operational units**
  - **Train staffs on best practices**

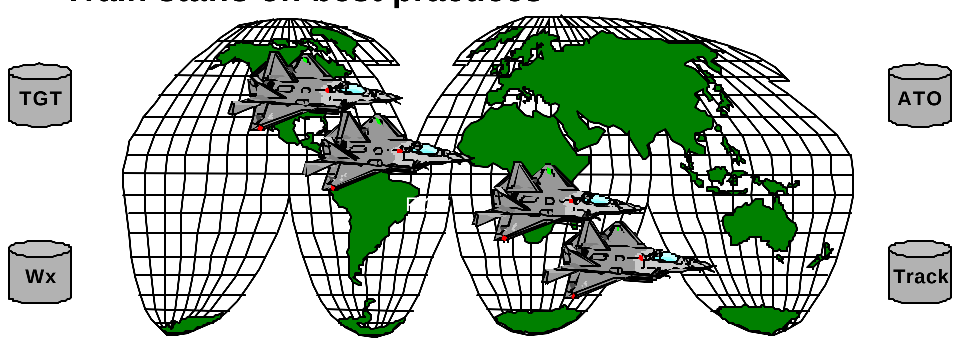

As the Air Force continues the process of evolving its Command and Control systems to meet the changing needs of the warfighters, it is time to rethink the role that data and information play in meeting these needs. Data is clearly recognized as critical to our Command and Control systems, since it is the data that tells us what targets to hit, the effect of our missions, the status of our forces, the readiness of our weapon systems and all the other information necessary to carry out combat missions. But data is often thought of as integral to a weapon system or to a command and control system. Or it is thought of as the property of an organization or a command. We want to make the case that data should be thought of as a strategic asset that belongs to the enterprise, and that it should be managed as an enterprise asset. Not every piece of data needs to be managed as an enterprise asset, but the Air Force needs a process for establishing the manner in which data can be sorted and its importance determined.

Data is a shared resource. When more than one organization/system needs the same data or information, there is no reason why there should be duplicate efforts to collect or to maintain that data. Yet we found numerous examples of duplicated maintenance of near-identical data. The use of different data management systems with different file structures may be a reason that justifies such duplicate maintenance - the Oracle system, for example, is used in TBMCS, while Sybase is used in AODB. The need to augment a database with additional fields that are for local use is another reason for duplication of the original database. Performance is another, especially if there are real-time requirements and network delays or the risk of network outage that would compromise the mission. These are all valid concerns, but there are technical solutions that can address each of them. There are real costs to not sharing that can be measured in both dollars and mission impact. The first problems to be addressed are the need to share data and the need for a commitment to do so.

Someone must be responsible for all data, and it must be possible to trace its pedigree. One of the purposes of sharing is to reduce the maintenance cost of multiple data copies. But this means that each of the data users must have confidence in it. It also means that the organization that maintains the data must know the different ways it will be used and will accommodate these users. This tends to place an additional burden on those who maintain the data. This will normally not be done with our existing systems since there is little incentive for it and it will, in fact, probably raise costs for the maintaining organization, even though it benefits the enterprise. The benefits of having a single responsible source are not only lower costs to the enterprise, but greater consistency in the data and a single authoritative data source. But this also means the maintaining organization will be responsible for meeting the needs of the entire user community in a timely manner. Today, there are no processes in place to insure that this happens.

Data exists at many levels and in many forms within an enterprise (e.g. in legacy systems, databases, metadata descriptions, images, sensor systems, etc.). Legacy systems frequently have data that is embedded in the code or is in some way usable only by the existing system. As we move toward a state where data is shared across the enterprise, both the operational data and the descriptive metadata must be shared. And we will be creating new knowledge by the fusion of information from multiple sources. Much of this will need to be shared as well. Many of these sharing issues were addressed in the JBI study through the use of functions such as "publish" and "subscribe." While those mechanisms provide a framework for sharing, further work is needed to establish specific processes to ensure accountability for data in the C2 enterprise.

Data must be protected. As there is greater sharing, particularly with coalition forces, policies and mechanisms must be provided to find the right balance between total protection and the sharing necessitated by the mission. This continues to be a major challenge for the Air Force.

Slide 42: Enterprise-oriented training

Train for the C2 Data Enterprise

- **C2 Operators must understand database complexity and its implications for operations**
  - **Analogy to teaching pilots about their airplanes**
    - **Must know enough to recognize problems and probable causes**
    - **Don't need to learn same level of detail as aircraft maintainers**
  - **C2 Operators must understand their data systems**
    - **Know when to trust data and when to confirm**
- **C2 Operators must be trained to share data**
  - **To share information from many sources**
  - **To contribute their expertise for the common good of the enterprise**
- **Also train the Data Managers, the System Developers, the Testers, etc.**

A crucial aspect of successful migration is the training of the C2 data operators. This training is distinct from that of the data managers, system developers and testers. Given the ubiquity of data in C2, every consumer and supplier of data in the enterprise will need basic training. The basic training should not be on database construction or the other aspects of database management- that should remain the province of maintainers and designers. Instead, the focus should be on ensuring that the members of the enterprise can access, understand, verify, supplement, and contribute to the C2 databases. This type of training does not currently exist in industry, academia, or the military and the USAF should pioneer a suitable program.

The proposed C2 database training has two objectives. The first is to provide operators with a high-level procedural understanding of the system sufficient to enable them to effectively extract information from the C2 databases. A key outcome of this high-level understanding is *confident decision-making*. In particular the operators must know when to trust the data and when and how to seek confirmation through other sources. This is analogous to the high-level training a pilot receives about how their aircraft work: they must understand the planes well enough to know when something is wrong and they must know the appropriate response. Their understanding is not the detailed understanding of an aircraft maintainer or designer, but rather is an operator-centric overview of the aircraft. As with aircraft, operators need an operator-centric view of databases. Unlike aircraft, where pilots are a subset of the general USAF population, almost everyone in the enterprise is a data operator.

As part of the high-level understanding, training must teach operators to *interact* with the C2 databases, rather than simply accepting the output. Confident decision-making assumes that the operator rejects the notion that "the computer is always right" and can identify possible errors in the data. This can lead to decision paralysis unless operators know how to get timely access to alternative sources of data in order to verify errors and supplement suspect data with better information. Since the C2 databases will continue to be distributed and multiply, training must actively promote the seeking of information rather than strict procedures.

The second objective of database training is the training of the C2 operator in *data sharing*. This objective is based on this study's view of C2 databases: a truly useful database system is constantly being updated and expanded in real time. As soon as an operator acquires a piece of information, it must be available to the larger enterprise. Since operators work in a distributed environment, they may not be good judges of the requirements of another unseen operator. Consider the case of targeting a bridge. The pilot has access to local targeting information but may not see that an unscheduled (or late) train is approaching. Another operator may observe that a train is currently running on that line, but yet be unaware of the pilot's mission to blow up the bridge. The information about the train must be made available and integrated with the pilot's understanding of the situation in order to truly meet the mission objectives.

Effective data sharing is unlikely to occur without specific training. Mechanisms for sharing data are not intuitive and are still primitive, so operators must (for the time being) make an additional effort. The distributed nature of the C2 enterprise may be such that there is no immediate reward for sharing data; for example, the external observer may never know that the information on the train running late prevented the pilot from blowing up the bridge and the train. Training will help the data operators contribute their expertise for the common good of the enterprise.

Slide 43: Migration Models Summary

Migration Models Summary

- USAF can create an effective development process to encourage enterprise data sharing and an effective data migration path
- AF needs a C2 acquisition approach which allows evolutionary migration
- A corporate-level entity is needed to encourage and manage the date migration process

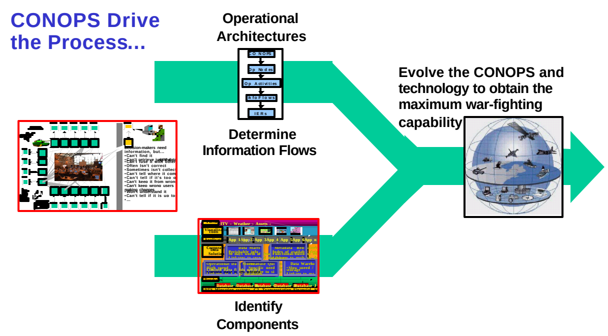

We believe there is a strong correlation between the migration approaches that have been used successfully in the commercial sector and those the Air Force needs. The commercial sector approach is governed primarily by such business considerations as “time to market” for products and a general need for financially sound decisions in the provision of the support infrastructure. The Air Force approach is driven by operational needs, but it is heavily constrained by acquisition policy, congressional appropriations (including restrictions based on the “color of money”) and organizational considerations (e.g., local decisions may not be in the best interest of the enterprise). If we are to apply lessons learned from the commercial sector, we need to address some of the institutional barriers existing in the Air Force. When we use the phrase “equip the C2 data enterprise,” we are talking about creating an environment where the Air Force can effectively apply migration lessons learned from the commercial sector.

- Policy—The Air Force needs to establish a policy for managing C2 data as an enterprise asset. This policy should recognize the notion that all data should be a communal or “share” asset, that someone should be responsible for that data, and that the use of data is a dynamic process with communities of interest expanding and contracting over time, just as applications also change dynamically. The key to the implementation of such a policy will be the establishment of processes for communities to define appropriate protocols and standards. A web-based sharing model should be considered.
- Evolutionary migration - Databases must evolve continually throughout their existence. This evolution should be a consequence of operational requirements, changing technologies, or changing communities of users. The databases need to be able to migrate without consideration for the systems

that use them. The whole process of migration must be thought of as a resource-constrained business decision where tradeoffs must be made between competing requirements. Note that we are emphasizing the migration of data, not the underlying command and control systems.

- Funding—Today's information systems (including the supporting databases) are generally funded as major acquisitions; the purchase and installation phase is followed by a sustainment phase. Since we are making the case for continuous evolution of the databases, there must also be a corresponding change in the funding model. We recommend that a "level-funding" model be adopted. We assume that the databases will be continuously updated and migrated as long as they remain in use, and that the user must have the flexibility to do so based on the operational need at hand.
- Managing the Process—The recommended strategy for dealing with databases at the enterprise level is significantly different from that followed today, and new processes are needed to implement such a policy. First, when dealing at the enterprise level, there are many tasks that, while they may be more expensive to an individual organization, will nevertheless benefit a larger community of users. This will require a cultural change in the way organizations make investments. There is a requirement for information sharing, and that implies the adoption of common protocols and standards. Prior attempts to adopt such standards have been too rigid and have generally failed. We believe that it will be necessary to exploit emerging and yet-to-be-invented technologies and that one must consider developing standards that apply to a "community of interest," rather than universally.

We have spent significant time and devoted much effort to thinking through these processes. We believe that such a process is necessary if the Air Force is to manage its data at the enterprise level. We have used the term "guiding hand" to describe an organization that will manage and oversee the migration process. This organization should be under the responsibility of the Air Force Chief Information Officer. Such an office would both provide guidance for those organizations that were responsible for managing enterprise-wide data and enforce the migration policies. We describe how this can work and propose a new C2 DB construct in the next, and final, section of this report.

Slide 44: Recommendations

Outline

- USAF C2 DB Findings
- Migration Success Factors
- Models for Successful AF migration
- **Recommendations**

In this, the final section of our report, we develop actionable recommendations to describe how the USAF can move toward a solution to the database migration problems we studied.

Slide 45: What is Needed

What is needed

- Corporate DB system evolution has been accomplished without significant new funding using evolutionary migration strategies
  - Learn from migration best practices
    - Save by freezing legacy functionality
    - Savings reinvested later in new functionality
- AF needs a C2 DB acquisition approach which allows evolutionary migration
  - Normal approach has been "big bang:" large capital investments, followed by sustainment phases -- and the cycle repeats
  - Data migration and integration processes should be level funded
    - Need to allow tradeoff and cost-reducing business decisions at every level

The USAF must drastically change its approach to the development and deployment of data-related C2 systems and databases. We can no longer afford the costs in either dollars or manpower that the current systems' development practices impose. The AF must level the funding of database migration and change from the process of buying new systems to a process of constant system evolution. This will lead us toward the sort of information management infrastructure described in the JBI and C2 reports discussed previously. In the remainder of this section we show how this need for evolution can be met, and we describe a new construct that should be put into place to allow this evolutionary migration.

Slide 46: How To Do It

How to do it

- **Establish and fund a "guiding hand" program office to promote migration of individual Air Force C2 systems to a component-based enterprise-spanning architecture (in conformance with JBI concepts).**
  - **Treat C2 data as an AF enterprise asset**
  - **Enforce AF CIO migration policy**
  - **Fund enterprise migration initiatives in individual programs (e.g. TBMCS, GDSS, N/UWSS, GCSS-AF)**
  - **Assist MAJCOMs and functional areas in migration strategies and selection of appropriate products**
  - **Coordinate with other appropriate DoD organizations**

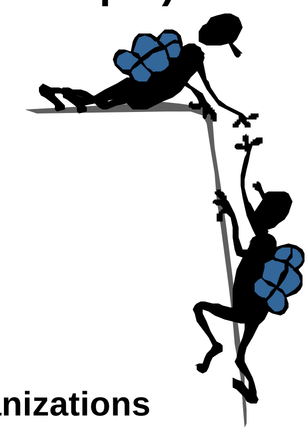

The notion of an office that exercises the role of "guiding hand" is the creation of an organization that provides consistent influence over the AF C2 database migration process. This organization needs explicit direction for its activities and funding/authority to exercise its mission. We note that recognition of the need for activities similar in nature to those previously discussed has occurred in the past. The "Link 16" program provides an example of an appropriate and realistic management structure that can be modified to serve this database problem.

First, the duties of a C2 Database SPO must be defined (we note some reservation in using the term SPO as the standard SPO mission may need to be somewhat altered in terms of management structure and tasking, as described in the next few slides). The C2 Database SPO will need to work with the Air Force Chief Architect's Office's OV-7 process to determine which information flow is the most critical and operationally relevant enterprise. Critical information flows will receive early attention to establish protocols and information configuration management processes which will scale across the Air Force-wide enterprise. A close association with the CIO's architecture office (AF/BIM) and the CIO office will therefore be necessary. The C2 Database SPO will also need to scour the civil commercial world to evaluate existing products and adopt and provide tools that will most suitably provide the efficient and cost-effective management of integrated databases. These tools must be made available, along with training for the contributing critical C2 nodes in the Air Force. The C2 Database SPO will also need to task the contributing components to both incorporate database constructs and accept information management responsibilities that align the databases for current and future evolution of the desired enterprise-wide database integration.

Second, the funding, requirements management and enterprise authority streams need to be defined. The following slides address these concerns.

Slide 47: Proposed Database Construct

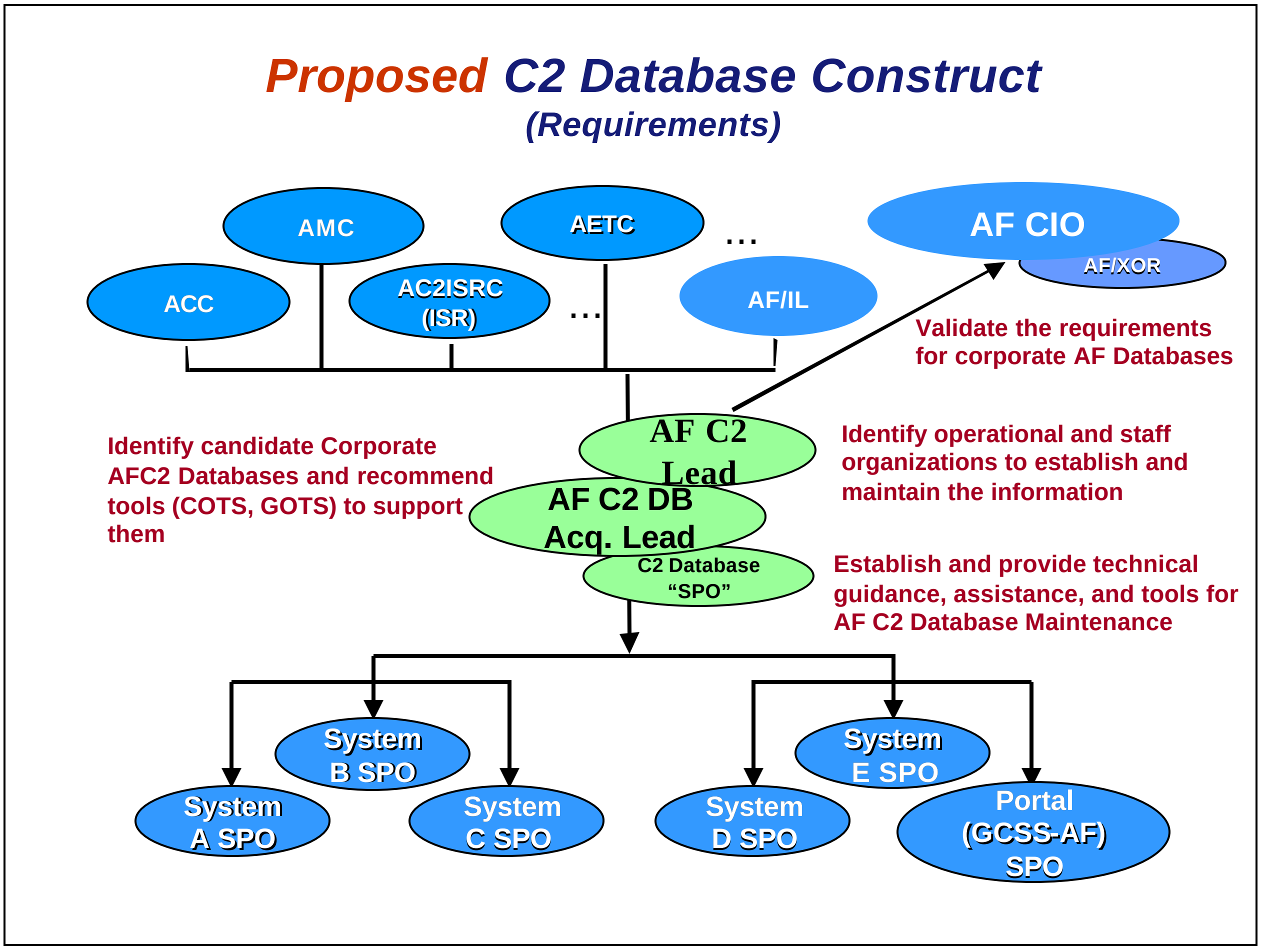

As was proposed for "Link 16" requirements, an AF C2 lead must be inserted as the organization responsible for managing interoperability. In the case of C2 databases, the enterprise acquisition authority (currently residing at ESC/CC) and the AF C2 lead (assumed to be the new DCS/XI) will identify candidate corporate AF C2 information needs in cooperation with the MAJCOMs and SPOs from developing weapons and other systems.

The C2 lead will assess the requirements for information flow that should be conducted on an enterprise level rather than an individual system level. A validation process will be conducted in concert with the AF/CIO office. Once validated, the AF C2 lead will recommend the operational and staff organizations responsible for populating the databases with information and for maintaining that information, and the AF C2 DB Acquisition lead (via the SPO) will recommend systems to be maintained and will publish this and similar information.

The validated requirement to manage AF corporate enterprise information will be passed to the executing SPOs. One of these SPOs will be the C2 Database SPO, who will provide enabling tools, implement the configuration control and information management coordination procedures and assess the integration progress achieved. Other SPOs will be the individual component SPOs responsible for the generation and ownership of information in the databases.

Slide 48: Funding Flow

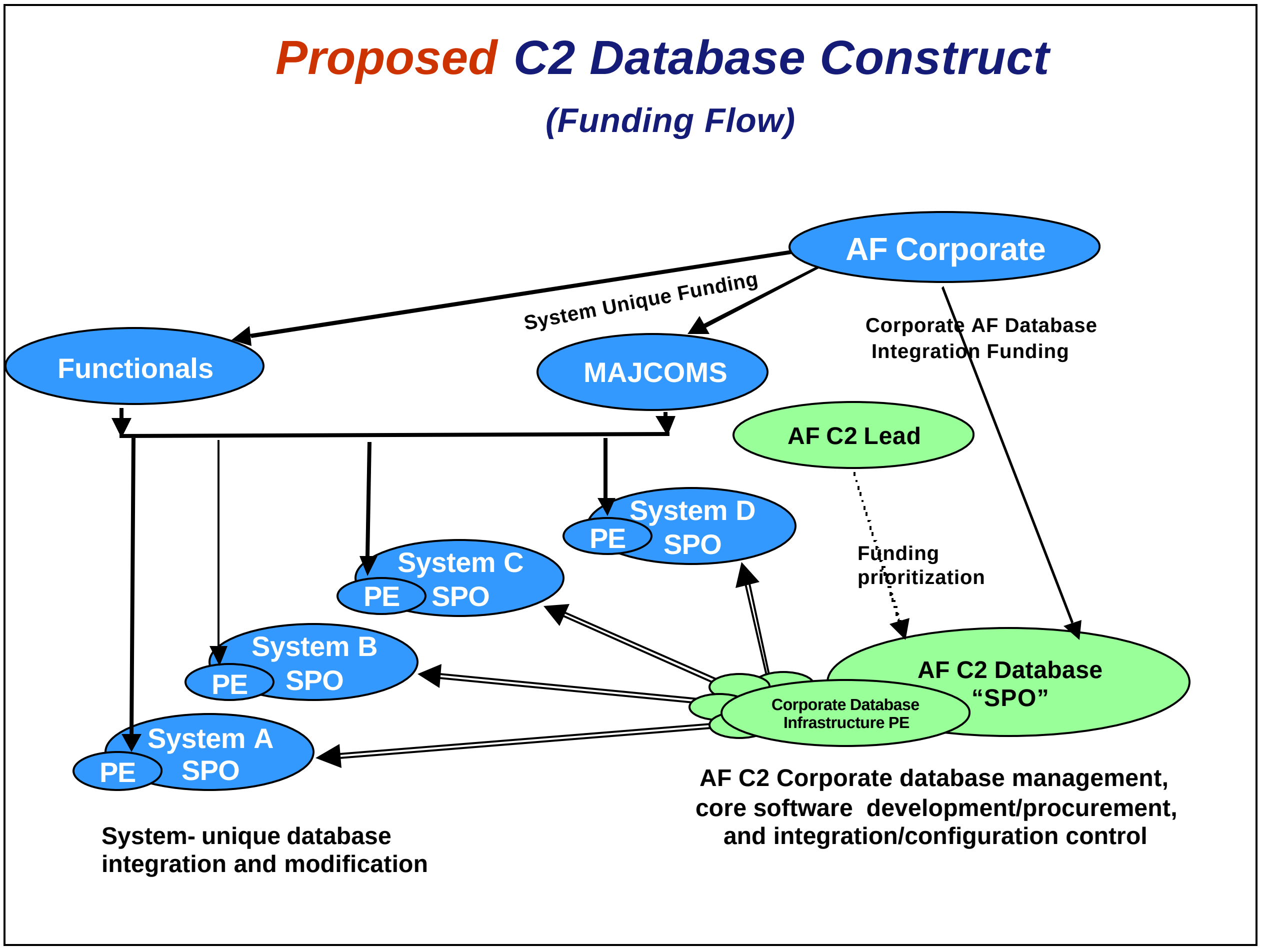

One of the most critical issues will be the flow of funds to enable the operation of the "guiding hand" process. Fund control provides the necessary authority to ensure long-term compliance and motivation among the responsible organizations to achieve the ultimate objectives of information integration.

During the programming phase, the proposed C2 Database Construct would require each SPO involved in the development of a system that owned or managed validated AF enterprise data to evaluate the challenges they must master before their system will conform to the integrated database solution. They would be tasked to assess the budget required to sponsor that effort. Then the AF C2 Lead organization would be required to validate the effort and funding estimates and to recommend to Corporate AF the gathering of the appropriate funds involved for delivery to the AF C2 Database SPO.

The normal system-unique funding would be provided via the appropriate Functional or MAJCOM sponsorship to the system SPO. Database integration funds would, however, be re-routed by AF Corporate to the C2 Database SPO. With the funds delivered from AF Corporate, the C2 Database SPO would manage the overall activities associated with the accomplishment of the C2 Database Integration. This would involve the tasking and resourcing of individual system SPOs for the incorporation of processes and technologies to achieve database integration. It would also involve the development or acquisition of tools and technologies to facilitate the management of the C2 Database Integration.

Such tools and technologies would be developed and maintained with the purpose of facilitating corporate achievement of the integrated database. They would be made available to the SPOs and the SPOs would be encouraged to use them since they are effectively "free" to the SPOs, but SPOs would also be permitted to build their own vehicles for information sharing so long as they abide by the protocols established for handling enterprise database information. The process will be designed to ensure the integration database SPO offers pragmatic solutions to the individual system SPOs by forcing the C2 Database SPO to be cognizant of the technical and operational challenges associated with the integration objectives.

Slide 49: Enterprise Management

Proposed C2 Database Construct

Enterprise Management

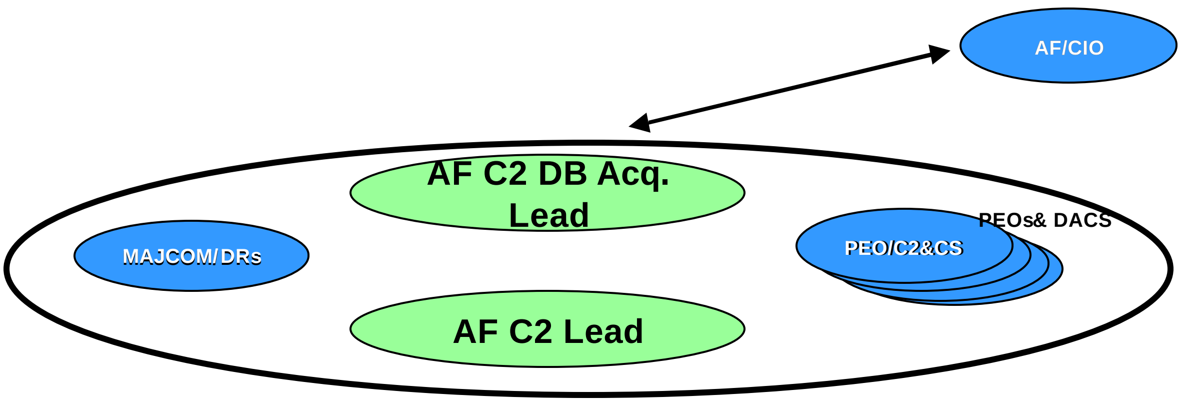

- **C2 Data Base General Officer Steering Group oversees C2 Database Enterprise Management**
  - **Report to AF/CIO**
    - **Recommend process correction/improvement**
    - **Suggest CIO policy direction**
      - Including removal of outdated policies
  - **Resolve budget/management issues**
    - **Identify remediation strategies**
    - **Resource actions exceeding C2DB PEO Scope**

A General Officer Steering Group (GOSG) will resolve budget and management issues that are not resolvable by the AF C2 Lead organization. Progress will be reported to another C2 Database General Officer Steering Group consisting of the MAJCOM/DRs, the PEOs/DACs, the C2 DB Acquisition Lead (the Integrated Enterprise authority) and the AF C2 Lead. This steering group will report to the AF/CIO, recommend any corrective actions towards the AF architecture process, and suggest CIO policy actions that would ensure adherence to the vision of integrated C2 data. Both the AF C2 Lead organization and the C2 Database SPO would provide support to the GOSG.

Under normal enterprise management procedures, the PEOs/DACs provide program direction to the SPOs, and warfighters evaluate requirements and development progress against those requirements. The AF C2 lead organization works in conjunction with those warfighter requirements and the C2 Database SPO activities to identify prioritization and solutions to issues that surface for development of the integrated database solution.

However, it is anticipated there may be disconnects between existing unique system designs and the need to satisfy distributed database capabilities. The cost and effort necessary for internal modification to systems may exceed the changes implied by simple adoption of database constructs. There may be additional internal communications, user interface, or application functionality modifications that are unintended consequences of the achievement of the integrated database. The problems of dealing with these issues, identifying remedial strategies, and resourcing the remedial actions may exceed the scope of responsibilities and the powers of the C2 Database SPO. A GOSG should therefore be made available to ensure proper perspective is preserved on the issues and an expedient and AF corporate-relevant solution is established. Anticipated tasks for the C2 Database SPO would be:

- *Identify the critical information flow.* This task will be conducted in concert with the AF/CIO and AF/Chief Architect's Office to determine which information flow activities involve enterprise relevant information. The determination that certain pieces of information constitute enterprise-relevant data will trigger compliance requirements on the data-owning component to induce conformity in the processes for C2 Database integration.
- *Develop/procure supporting tools*. The C2 Database SPO will be responsible for the location and procurement of tools that will help the database integration process in the civilian market. These might include such items or processes as hypertext management tools, tools for building data objects and configuration management tools.
- *Task contributing SPOs and manage the configuration control processes necessary to assess progress*. While data designated as "enterprise-relevant" will be required to follow the C2 database integration processes, it is expected that the recommended processes will often be suitable for data within C2 components. An adherence assessment will be conducted upon each contributing component SPO to assess their progress not only towards the specific critical information flow data element integration process, but also their progress towards a systemic adoption of the principles of data integration.

Slide 50: Recommendations Details

Recommendations Details

- **Direct the CIO to institute and manage a process to continuously evolve USAF C2 enterprise data structures (and the C2 systems they support) and to insure high quality data for timely decision-making. (SECAF, CSAF)**
  - **Develop the USAF C2 Enterprise Operational Architecture and identify the AF's IERs(PDAS/BIM working groups a good start)**
  - **Incrementally upgrade functionality, drawing from COTS and research products**
- **Direct the MAJCOMs and functional areas to manage data in accordance with the established process**

The mission of the organizational structure described in the previous slides must be managed at the highest levels. We believe the USAF must set up the process previously described, and it must direct the CIO to manage this process of continuous evolution of AF C2 databases and systems. The overall process would help the operational needs of time-critical missions to be recognized and dealt with across the data enterprise. In addition, the MAJCOMs and functional areas must be directed to comply with this approach and to assign people to the General Officer Steering Groups described above.

Slide 51: Other Recommendations

Other Recommendations for C2 Data Migration

- **Permit local initiatives**
  - **Allow standard databases to be locally augmented/expanded**
  - **Provide consulting/help desk services to unit-level development efforts**
- **Prioritize migration efforts according to cost reduction/mission risk tradeoff**
- **Mandate and enforce DoD XML tag use and registration in the AF**
  - **CIO architectural working groups a good start**
- **Monitor and evaluate the key features of DB migration efforts**
  - **Develop a (web-enabled, on-line, marked-up) repository of lessons learned**
- **C2TIG must include training on underlying principles of IM systems, and the implications thereof on operations**
- **On an ongoing basis, review information system policy directives and remove those that are unwieldy, unrealistic, or out of date**
- **Assurance mechanisms need to be planned and built into AF DB systems to support the needs and possible uses by a variety of users (esp. coalition)**
- **Direct research investment in several strategic areas**
  - **Information access and control, Service-based middleware, Semantic interoperability**

In addition to the recommendations mentioned above, there are a number of other recommendations we believe are important to the long-term success of this endeavor. We summarized these above and discuss them in greater detail below:

- *Permit local initiatives.* It is clear that as time passes, the airmen and women of the using units will gain both familiarity with information technology and availability of COTS tools with which they can better pursue their needs. Policies and procedures should be instated to allow these "business units" to develop individual competencies, but in a manner as consistent as possible with overall enterprise needs. We believe two particular approaches are important:
  - *Allow standard databases to be locally augmented/expanded.*
  - *Providing consulting/help desk services to unit-level development efforts should be a high priority*. The first of these would prevent many of the current ad hoc interfaces by simply creating more consistent means by which the units that have the ground truth can "annotate" the values in system databases. For example, the unit that can measure the length of a runway should be able to enter the fact that it is different from the posted value, without this requiring a weeks-long process or being contrary to policy. These local values need not be considered definitive, but where they differ with a value in the system of record, that fact should be noted electronically and high priority should be given to the work of rectifying these differences if they begin to take on an operational significance (for example, a fighter is to be landed on that runway). Secondly, the AF C2 Database Acquisition Lead (currently ESC) should create means to help the units to

develop their own systems in a way that is compatible with overall AF needs. We suggest the nascent UDA MAN program at ESC be examined as a possible starting place for this development.

- *Prioritize migration efforts according to cost reduction/mission risk tradeoff.* As we have described previously, a process must be put in place to determine the systems that can be taken offline so that the maintenance costs of these systems can be reclaimed for other use. This does not require waiting for the organizational structures we described previously, and we suggest each MAJCOM and functional area be asked to start identifying such systems immediately. In addition, current efforts to develop an overall AF C2 CONOPS are critical to the prioritization of systems for modernization, and we suggest those efforts be even more focused on actual IERs than is currently the case.
- *Mandate and enforce DoD XML tag use and registration in the Air Force*. As discussed in the technical section, XML is a useful technology for AF data integration (even if not the magic bullet some people believe it will be). We recommend that the USAF work out means to start developing critical tag sets that will be of use in this migration (without going overboard and trying to create a single schema that all must follow - a disastrous course of action). The means to accomplish this is to allow subgroups with critical needs to work out high-level tags that they wish to see within their functional missions. The CIO has created some working groups to start this process, and we recommend that these be monitored and adjusted as needed, and that they continue in an ongoing process that helps to create important starting places for the data migrations that will follow.
- *Monitor and evaluate the key features of DB migration efforts*. Put simply, the Air Force needs to make it easier for best practices to be found and shared. Recognition should be available for those units that are doing the best job of this, and these best practices should be shared. We suggest that the Air Force develop a repository of lessons learned (and that repository be available continually as a web-enabled online resource marked up with XML or other advanced tagging to make the searching of the archive easier).
- *C2TIG must include training on underlying principles of IM systems, and the implications for operations*. This was discussed in detail on page 116, and we simply recommend this be put in place at the earliest possible time.
- On an ongoing basis, review information system policy directives and remove those that are unwieldy, unrealistic, or out of date.
- A necessary component of the CIO's involvement and the General Officer's Steering Group focus on this problem will be to, on an ongoing basis, make sure that the policies which the Air Force uses to control and standardize data interoperability do not become prohibitively restrictive or stifle best practices. In addition, requirements coming from above (from JROCs and other organizations) can conflict with Air Force directives in a way that makes obedience to all policies impossible and encourages the (costly) solution of applying for waivers. The CIO process can both examine Air Force policies that are outdated and advocate the abolition of prohibitively expensive external policies.
- Assurance mechanisms must be planned and built into Air Force database systems to support the needs and possible uses by a variety of users (especially coalition partners).
- *This report, like most other USAF reports, has been forced to ignore some important security issues*. The reason behind this is relatively simple: security concerns encompass the whole body of Air Force and DoD systems, and any review of security that covers only a single system will be incomplete. Any security mechanisms that were peculiar to databases would be incomplete, and would be "an application of Band-Aids where stitches were required." We recommend strongly that the acquisition leads for Air Force systems start paying considerably more attention to the overall needs of those systems, rather than trusting someone else (for example DISA's SIPRnet) to provide the needed security. We believe the experimental spirals described in last year's report (and mentioned on page 110) are the right places to explore the issues, to model possible hostile approaches and to determine where specific policies and procedures (technical and cultural) can be used to provide critical additional security to systems whose compromise could put Air Force personnel and operations at risk.

- *Direct research investment in several strategic areas*. Our focus in this report has been "tactical" - we have been describing processes and procedures that can and must go into effect *now*, to start the immediate improvement on a problem that will get worse as time goes on. However, at a "strategic" level, we also found the Air Force to be under-funded and possibly lagging behind the Army and the Navy. DARPA has been funding a number of critical technologies, and the AFRL needs to be tasked (and funded) to work with the research community to help steer the developments of these technologies and to transition them to Air Force use. Three critical areas of investigation are:
    - Information access and control.
    - Service-based middleware.
    - Semantic interoperability.

All three are critical for dealing with the long-term information management needs of Air Force C2 as described in previous AFSAB studies. All three of these fall under the aegis of the current Joint Battlespace Infosphere (JBI) research under way at AFRL/IF, and we repeat a recommendation made earlier in this report (page 112) - the JBI work is critical to the future data/information needs of the Air Force and needs more funding both at the labs and in the process of testing these technologies for operational needs (for example in CAOC-X and JEFX activities).

Slide 52: Conclusion

Conclusion

C2 data management is a critical mission function spanning MAJCOM boundaries.

AF leadership must recognize the value of C2 data enterprise management and take significant steps to achieve the Air Force Vision.

You Can Get There From Here

At the conclusion of our section on the current Air Force practices, we made it clear that without organizational change the Air Force simply cannot overcome its current DB problems in the C2 arena. Despite a few promising directions, the policies and processes we use (which buy information systems as if they were airplanes) cannot be successful at providing the infrastructure necessary for the global vigilance, reach and power which constitute the future vision of the US Air Force. We have described in this report a set of organizational, operational, and technical steps that can be taken to change matters. An understanding of the importance of C2 Data Management to the Air Force can allow us to take the critical steps needed to overcome this problem and achieve our vision.

(This Page Intentionally Left Blank)

Appendix A

Database Migration Described

A.1 The Definition of Migration

We define migration as movement to a new system or subsystem. In Chapter 5, we have shown that migration to a new paradigm is desirable. The need for migration is, however, independent of the paradigm chosen for the future.

A.2 Benefits and Liabilities of Migration

A.2.1 Enabling New Applications and Technologies

The principal benefit of migration is a fresh start, but even the most successful migration carries the corresponding liability of a temporary loss in flexibility due to the loss of the experience associated with the existing system. Offsetting that loss is another important benefit, the enabling of new technologies and applications available in the new database, a benefit whose apparent value is magnified by the tendency of systems to deliver fewer of their starting benefits as they age.

A.2.2 Providing For New Data

New applications typically require new data. The process of changing a system to deal with a broader range of data can, if it is not carefully planned, require changes that cost more than a system replacement. If the data already exists in another system, bridges must be built to move the data across to the new applications. At times that data may not be of adequate quality for the new application. In such cases a parallel effort is needed to alter operations in the source system. The existence of such bridges complicates system management, to the extent that no systems can be upgraded since the extent of related changes becomes unpredictable.

A.2.3 Breadth

Even though our objective is labeled "database migration," it must be recognized that databases are just one component of the system and that there are several components we must consider. Some are commonly recognized as parts of databases, and the importance of others may not be obvious until the migration begins. The following are some of the concerns we will discuss below:

- Subsection A.3: The data contained in the databases.
- Subsection A.4: The metadata that describe the data, both to users and programs.
- Subsection A.5: Business rules that constrain the data or its processing.
- Subsection A.6: Application programs that produce results from the databases.
- Subsection A.7: The human aspects of migrating to new systems.

A.3 Actual data migration requirements

Moving data values stored on the system files from a legacy system to a modern system may be a relatively simple process, if the quantity and distribution are modest and the data are clean. Problems arise mainly from data errors that become visible when the tighter controls of a new system impose themselves. For instance, if an objective assessment, such as "suitability" was a textual field in the original application but is coded in the newer one (to provide, perhaps, for further processing), then unusual terms will have to be manually recoded.

While relational database management systems use a common model and generally provide tools for downloading and uploading the entire contents into a neutral format, restructuring will still be needed if the schema changes, as is likely for a migration that involves maturation of the architectural concepts. Even more restructuring of data is needed when the conceptual model changes; an example of this sort of change would be a movement from a relational base to an object-oriented or hierarchical structure. Movement into an XML world often implies a hierarchical restructuring, although that is not strictly necessary.

While the restructuring is typically only modestly complex, the large quantity of data that must be transformed will always pose some problems. Of particular importance are data errors that were not recognized or were ignored in the legacy structure. These may cause errors and even massive failures during such a migration. It is important to remember that data do not live in isolation. While data can often be converted and shipped relatively simply, we must also migrate related information so it can be understood.

A.4 Metadata migration

Data dictionaries and standards (such as those being developed for DoD) collect some of the metadata, but rarely to the extent needed for automation. The form that metadata take differs greatly among systems. A metadata often describes the data, preferably in such a way that both people and programs can manipulate it correctly. A great deal of manual work is required. This area is one where great progress is being made, so much so that general recommendations are hard to establish. The upcoming XML standard provides a basis for sharable metadata, but it does not cover all the facets that are needed to describe data for reliable processing.

Unfortunately, in most older systems metadata was only defined at design time, and no source information has been kept or maintained. Analysis and study are required to recover the intent of the data and the constraints being imposed. Since the volume of metadata is small, manual metadata migration is feasible, but a good deal of expertise is needed to accomplish a successful metadata migration.

A.5 Business rules migration

An important subset of metadata, best treated distinctly, are business rules. A simple rule might be

Car rental cost = daily charge + insurance + late charge + return cost

Business rules define the constraints imposed by data during processing but cannot be imposed on the static storage of the data, since similar data may undergo many different permutations. For instance, if a business rule states:

New hires cannot by paid more than $30K

It is not feasible to limit the general salary field in the database to $30,000.

The rules and constraints that determine what data can be updated by whom, what updates have to be reflected in other portions of the information systems, what events have to be reported, and how and to whom information is to be disseminated are built into the code of application programs. Other rules will control and guide the control of activities flow (also known as "the work flow") in major systems.

In most modern systems, business rules are not explicit but are embodied in programs. Many of these rules require changes as settings change, new types of data are obtained and new functions are added to the systems. The process of changing the code of the application programs is tedious and error-filled. Testing of updated application programs must be rigorous, because the changes made can have unforeseen effects.

By extracting these rules, coding them into clean high-level applications and automating the interpretation of such rules, future program managers can attain substantially lower maintenance costs. However, modern programming approaches make it desirable that the rules become explicit, making the processing easier to understand and to modify when the situation warrants.

When business rules are incorporated in programs, the process of their specification, implementation and validation is likely to be lengthy, because many rules are combined in a program and they are often deeply integrated to achieve high performance. The execution of business rules becomes costly when they relate two items that are remote in space or time. Charges in such circumstances are allowed only if the sum of prior charges plus the current charge is less than the budget.

Much research has been performed on the problems inherent in moving to explicit business rules; an example will be illustrative. For instance, when explicit business rules are implemented in object technology, they will generate smaller code sizes (they are also easier to change). Some implementations use AI technology derived from OPS5.

Database developments themselves have adapted rules, under the rubric Active Databases: Ode, HIPAC, ADAM, Sentinel, and Trigs. Their acceptance has been limited because the collection and execution of rules adds further to DBMS's complexity and creates a uniqueness that many customers want to avoid. In applications we can insert rules into a distinct layer, as in the Java Expert System Shell (JESS) with the Rete reasoning engine. It is possible to use domain classes as JESS classes. Such a mechanism is easily accessible to Java programmers, but the process of placing the rules into applications raises the issue of consistency. How can application-based rules guarantee DB correctness when multiple applications are permitted access to a database? Under such circumstances, all rule sets have to be consistent.

Rules, whether centralized or distributed, can be ambiguous and can create conflicts. It may be wise to always provide manual overrides, with a log, that allow subsequent repair. For instance, if an essential new hire requires more than $30,000, then the manager's decision should not be permanently disabled because of an extant rule. Such concerns may especially crucial when we must serve wartime priorities.

Actual rule languages are still difficult for hastily trained staff to use. Examples shown in brochures always show simple cases, but rules can quickly become very complex and a high level of expertise is needed for their construction and maintenance.

A.6 Applications migration

Databases serve applications. Although some simple operations can be carried out by the tools that come with a database system, more complex and situation-specific operations will require application programs. Included with any DBMS are procedures for updating, managing the storage, querying the contents and generating simple reports. Often quite extensive report generators, warehousing, and data-mining tools can be obtained from the same or competing vendors.

Task-specific applications often require specialized programs. In a pure migration setting, the prior versions of the system already support most functionalities so that prior applications programs exist. The extent of reuse of existing programs can differ greatly, and it varies with the type of migration. Some examples include the following:

- ***Migration from one relational database to a more modern version.*** In this case, all services provided to the prior applications should be available and the application migration can focus on the incompatibilities of formats and interfaces that will invariably arise.
- ***Migration from one database to a modern, relational database system.*** Many older databases provide services that require the writing of new and often complex SQL transactions in a relational environment. Old interfaces have to be excised, cleaned and rewritten. Sometimes mediation technologies (see Chapter 4) must be employed to resolve major incompatibilities, especially of the source databases, because they must now serve a plethora of applications. However, little design and analysis is needed, since at a high level the application functionality is to be replicated.

- ***Migration from a file-based storage system to a modern, relational database system.*** Many older applications were written to use dedicated files. Often some of these files were generated by other, older systems. Movement to shared databases reduces the redundancy, inconsistency and obsolescence inherent in these legacy approaches. The file access and update code of these applications must be excised and rewritten. Before investing in such efforts, these applications should be carefully studied. Some may not be needed at all, because better paths to make the information available exist and standard tools can be used to replace all essential functionality. It is a poor practice to simply make the specifications for the migrated version of the application identical to the old specifications, because often the reasons for features of such legacy systems have long fallen by the wayside - or should have.
- ***No useful prior software is available.*** A problem in application migration arises when software has been written in obsolete languages. Even if the language was wonderful, if it is a language that is no longer widely taught, expending effort on a conversion that will be costly to maintain over the long term should be avoided. If the required concepts for an application are well understood, then writing software to order (often using simple tools as report generators and scripting languages) is a wise choice. We find that over-engineered software has fallen out of practice in commercial applications, since time-to-market is more important than performance or potential longevity

Any new and revised applications should be compatible with modern interface standards. Standards will keep changing, and their long-term viability cannot be guaranteed. When multiple standards are available, the standard that is currently more popular in practice should be preferred. Standards that derive their mandate merely from a governmental fiat are rarely chosen.

In traditional applications, much code was devoted to creating user-friendly interfaces, although the end-user may not perceive them as such. The dominance of the web and the ubiquity of browsers have made this application aspect nearly moot. By providing a simple HTML or XML interface much effort can be saved, and some of the reasons for salvaging applications can disappear entirely. This argument, when combined with the arguments presented above, can greatly simplify application migration.

A.7 People

People are a key aspect of migration. We need trained personnel who can carry out the job. We often hear that there is a shortage of database administrators. Furthermore, database administrators often change jobs. Therefore, we need to ensure that there is sufficient personnel support before we carry out a migration effort. We also need commitment and resources from management. Data ownership is yet another issue. We need to clearly define the ownership of the data and the roles and responsibilities of the involved personnel. For example, does the legacy administrator still have ownership of the data migrated to the new platform? Successful migration requires that we answer such difficult questions.

A.8 Alternatives

At times there is just not sufficient benefit to warrant the investment that a migration can entail. An isolated program in a stable setting that does not require much maintenance is probably best left alone. Any maintenance will be costly, since documentation is likely to be poor, programming expertise for legacy technology rare, and contracts expensive. To enable the use of legacy systems in a more modern distributed computing context, one must provide a communication interface to access data.

A.9 Summary

In any practical migration, information and programs in all these categories must migrate together, although the representation, complexity, and volume will differ greatly among them. Because of these differences we observe that the approaches taken differ as well.

For each migration task, there is a process for accomplishing it. And there should be a set of metrics associated with the processes. These metrics can include the system cost, development time, testing and validation time, distribution and synchronization, interoperability and retraining costs.

(This Page Intentionally Left Blank)

Appendix B

Terms of Reference

Migration of Data Bases for Command & Control

Terms of Reference

BACKGROUND

The rapid evolution to modern information systems creates a very major problem in the continuing viability of the legacy databases. The SAB just completed a major summer study on Command and Control, and the successful implementation is going to require that the many databases, which are used throughout the C2 enterprise, can be successfully migrated to emerging and future systems.

STUDY PRODUCTS

Briefing to SAF/AQ, ESC/CC, and C2ISRC/CC in July 2001

CHARTER

The study will review databases that are involved in command and control systems and processes, and make an assessment of the state of their accessibility by the emerging systems associated with TBMCS. The study can consider database issues such as standards, management practices, etc. as appropriate, but will accomplish the following:

- Make recommendations on the strategy, processes, and technical detail (if helpful) on assuring the continuing viability of the data contained in the legacy databases.
- Make recommendations on the further migration of the databases to a Joint Battlespace Infosphere environment over the longer term.

(This Page Intentionally Left Blank)

Appendix C
Study Team

Study Chairman
Prof. James A. Hendler

SAB Military Director
Lt Gen Stephen B. Plummer

Senior Civilian Advisors
Mr. John Gilligan, AF/CIO
Dr. Louis Metzger, AF/ST

General Officer Participant
Brig Gen Gil Hawk

SAB Executive Director
Col Gregory H. Bishop

SAB Study Executive Officer
Lt Col Paul Schubert

Panel Chairs
Operations Panel: Maj Gen (r) Eric Nelson
Migration Panel: Mr. Thomas Saunders
Commercial Panel: Prof. Gio Wiederhold
Viability Panel: Dr. Duane Adams

Operations Panel

Maj Gen (r) Eric Nelson, Chair
Aerospace Industry Consulting Services

Brig Gen Gil Hawk
U.S. Air Force

Mr. Robert Beaton
Private Consultant

Mr. Mike Livingston
BTG, Inc.

Prof. Robin Murphy
University of South Florida

Mr. Thomas Wade
TRW, Inc.

Mr. Gary Edwards
ISX Corporation

Mr. Thomas Andrew
Private Consultant

Migration Panel

Dr. Bhavani Thuraisingham, Co-Chair
MITRE Corporation

Mr. Thomas Saunders, Co-Chair
MITRE Corporation

Ms. Teresa Lunt
XEROX Corporation

Dr. Scott Renner
MITRE Corporation

Prof. V.S. Subrahmanian
University of Maryland

Prof. Alan Willsky
Massachusetts Institute of Technology

Viability Panel

Dr. Duane Adams
Carnegie Mellon University

Mr. Michael Brenton
Logicon Sterling Federal

Prof. Eugene Spafford
Purdue University

Prof. Nancy Leveson
Massachusetts Institute of Technology

Technology Panel

Prof. Gio Wiederhold
Stanford University

Mr. Thomas Clark
AFRL/IFSE

Mr. Scott Fouse
ISX Corporation

Appendix D
Acronyms and Abbreviations

3NF	Third Normal Form
5D	Demand Driven Direct Digital Dissemination
AAT	ATO/Airspace Control Order Tool
ABDM	Air Mobility Command Business Decision Model
ABP	Air Battle Planning/Plan
AC2ISRC	Airspace Command and Control, Intelligence, Surveillance and Reconnaissance Center
ACC	Air Combat Command
ACDB	Air Campaign Database
ACFP	Advanced Computer Flight Plan
ACO	Airspace Control Order
AD	Airspace Deconflict Tool
ADANS	AMC Deployment Analysis System
ADC	Air Defense Command
AF	Air Force
AFMC	Air Force Material Command
AFOTEC	Air Force Operational Test and Evaluation Center
AFRL/IF	Air Force Research Laboratory, Information Systems Directorate (also AF/IL)
AFSPACE	Air Force Space Command (also AFSPC and AFSPACECOM)
AFSOC	Air Force Special Operations Command
AFSOF	Air Force Special Operations Forces
AF/XI	Deputy Chief of Staff for Warfighting Integration
AF/XOR	Deputy Chief of Staff for Air and Space Operations, Operational Requirements
AIM	Airlift Mission Planning Tool
AIT	Automated Identification Technology
AMC	Air Mobility Command
AMOG	Air Mobility Operations Group
AOC	Air Operations Center
AODB	Air Operations Database
API	Application Program Interface
ASOP	Air Support Operations Prototype
ATO	Air Tasking Order
AUSD(TP)	Assistant Undersecretary of Defense for Technology Policy
B2B	Business to Business
B2C	Business to Customer
B2E	Business to Employee
BDA	Battle Damage Assessment
BDSS	Business Decision Support System
C2	Command and Control
C2IDD	Command and Control Interface Design Document
C2IPS	Command and Control Information Processing System
C4ISR	Command, Control, Communications, Computers, Intelligence, Surveillance and Reconnaissance

CAFMS	Computer-Assisted Force Management System
CAMPS	Consolidated Air Mobility Planning System
CAMS	Core Automated Maintenance System
CAOC-X	Combined Air Operations Center (Experimental)
CAPS	Consolidated Aerial Ports System
CAST	Close Air Support Tool
CCPDS-R	Command Center Processing and Display System
CDE	Corporate Data Environment
CDO	Corporate Data Office
CDS	Corporate Data Solution
CEDI	Commercial Electronic Data Interchange
CFM	CONUS Freight Management
CGI	Common Guard Interface
CINC	Commander in Chief
CIO	Chief Information Officer
CIS	Combat Intelligence System
CMARPS	Combined Mating and Ranging Planning System
CMAS	Cheyenne Mountain Air Station
CMOC	Cheyenne Mountain Operations Center
CMOS	Cargo Movement Operations System
CMU	Cheyenne Mountain Upgrade
CNA	Computer Network Attack
CND	Computer Network Defense
COE	Common Operating Environment
COMPES	Contingency Operations/Mobility Planning and Execution System
CONOPS	Concept of Operations
COTS	Commercial Off the Shelf
CPAF	Cost-Plus Award Fee
CPRP	CIO Program Review Panel
CRM	Customer Relations Management
C/S	Client-Server
CTAPS	Contingency Theater Automated Planning System
DAAS	Distributed Auto Archive Services
DAAS	Defense Automated Addressing System
DAC	Defense Acquisition Circular
DARPA	Defense Advanced Research Projects Agency
DB	Database
DBMS	Database Management System
DCAPES	Deliberate and Crisis Action Planning System
DDDS	Defense Data Dictionary System
DII COE	Defense Information Infrastructure Common Operating Environment
DISA	Defense Information Systems Agency
DoD	Department of Defense
DODAAC	Department of Defense Activity Address Code
DTD	Data Tag Definition
DTE	Defense Transportation Enterprise
DTJRT	Defense Transportation Joint Reference Table

DTRACS	Defense Transportation Report and Control System
DTS EA	Defense Transportation System Enterprise Architecture
DTTS	Defense Transportation Tracking System
DW	Data Warehouse
EA	Enterprise Architecture
EAF	Expeditionary Air Force
EDB	Enterprise Database
EMC	Execution Manager Control
EMR	Execution Manager Replanner
EMRDB	Execution Manager Replanner Database
ESC	Electronics Systems Center
ETL	Extracted, transformed and loaded
EUCOM	U.S. European Command
EW	Enterprise Workstation
FOC	Full Operational Capability
GATES	Global Air Transportation Execution System
GCCS	Global Command and Control System
GCCS-I^3	Global Command and Control System Integrated Imagery and Intelligence
GCCS-M	Global Command and Control System - Maritime
GCSS	Global Combat Support System
GDSS	Global Decision Support System
GEOLOC	Geographic Location Code
GFE	Government-furnished equipment
GMD	Ground-based mid-course missile defense (also Group mux and/or demux)
GOPAX	Groups Operational Passenger System
GOSG	General Officer Steering Group
GOTS	Government off-the-shelf
GTN	Global Transportation Network
GTN 21	Global Transportation Network for the 21st Century
GUI	Graphical User Interfaces
HMI	Human Machine Interface
HTML	Hyper-text Markup Language
IATA	International Air Transport Organization
IBS	Integrated Broadcast System
ICBM	Intercontinental Ballistic Missile
ICD	Interface Control Documents
ICE	Integrated Command, Control and Communication System
IDD	Interface Design Document
IDM	Intelligence Data Management
IER	Information Exchange Requirements
IM	Imagery Management
IMDD	Integrated Movement Data Display
IMT	Integrated Management Tool
INMARSAT	International Maritime Satellite
IOC	Intelligence Operations Center
IPL	Integrated Priority List
ISC2	Integrated Space C2

IT	Information Technology
ITO	Integrated Tasking Order
ITV	In-transit Visibility
JALIS	Joint Air Logistics Information System
JBI	Joint Battlespace Infosphere
JDP	Joint Defensive Planner
JESS	Java Expert System Shell
JEFX	Joint Experimental Force Exercise
JFACC	Joint Force Air Component Commander
JOPES	Joint Operational Planning and Execution System
JPT	JFACC Planning Tool
JROC	Joint Requirements Oversight Council
JRTLP	Joint Reference Table Logistics Project
JTT	Joint Targeting Toolkit
LDM	Logical Data Model
LMMS	Lockheed Martin Missile Systems
M2K	Mobility 2000
MAD	Mutually Assured Destruction
MAJCOM	Major Command
MCD	Migration Completion Date
MCCC	Mobile Command and Control Centers
MD	Missile Defense
MEPED	Military Equipment Parametrics Engineering Database
MIDB	Modernized Integrated Database
MILSTAMP	Military Standard Transportation and Movement Procedures
MOT&E	Mission Operational Test and Evaluation
MTMS	Maritime Tactical Message System
MW	Missile Warning
M/Y	Man-year
NASA	National Aeronautic and Space Administration
NAVSPACE	Naval Space Command
NBMC	NORAD Battle Management Center
NCA	National Command Authority
NORAD	North American Aerospace Defense Command
N/UWSS	NORAD/USSPACECOM Warfighter Support System
OA	Operational Architecture
ODS	Operational Data Store
O&M	Operations and Maintenance
OLAP	On-line data Analysis Programs
ORD	Operational Requirements Document
OS	Operating System
OT&E	Operational Test and Evaluation
PACOM	Pacific Command
PAFB	Peterson Air Force Base
PDM	Physical Data Model
PE	Program Element
PEO	Program Element Officer

PICCS	PACAF Interim Command and Control System
PKI	Public Key Infrastructure
PMO	Program Management Officer
R&D	Research and Design
RCAPS	Remote Consolidated Aerial Port Subsystems
RDF	Resource Description Framework
RF Tag	Radio Frequency Tag
S&T	Science and Technology
SAB	Scientific Advisory Board
SC	Communications Staff
SIL	System Integration Laboratory
SMM	System Maturity Matrix
SOE	Sequence of Events
SOO	Statement of Objectives
SPACECOM	U.S. Space Command
SPADOC	Space Defense Operations Center
SPLC	Standard Point Location Code
SPO	System Program Offices
SQL	Structured Query Language
SSG	Standard System Group
STRATCOM	U.S. Strategic Command
TACC	Tanker and Airlift Control Center
TACS	Theater Air Control System
TAP	Theater Air Planner
TAPDB	Theater Air Planner Database
TBMCS	Theater Battle Management Core System
TBMCS-UL	Theater Battle Management Core System – Unit Level
TC ACCIS	Transportation Coordinator Automated Command and Control
TCT DB	Time Critical Target Database
TDC	Theater Deployable Communications
TIM	Technical Interchange Meetings
TLDM	Transport Logical Data Model
TMDS	Table Management Distribution System
TPFDD	Time-Phased Force and Deployment Data
Track DB	Tracking Database
TRANSCOM	U.S. Transport Command (also USTC)
TRB	Technical Review Board
TSA	Target System Architecture
TWM	Targeting and Weaponeering Module
USAF	United States Air Force
USAFE	U.S. Air Forces Europe
USSPACECOM	U.S. Space Command
USTC	U.S. Transportation Command (also TRANSCOM)
VAMP	Vulnerability Assessment Management Program
VPN	Virtual Private Network
WAI	Web Application Interface
WCCS	Wing Command and Control System

WPS	Worldwide Port System
XML	Extensible Markup Language

Appendix E
Organizations Consulted

United States Transportation Command

North American Aerospace Defense Command

United States Space Command

United States Central Command

United States Special Operations Command

Air Combat Command

Pacific Air Forces

Ninth Air Force

Electronics Systems Center

Gunter Air Force Base

Hanscom Air Force Base

Defense Information Systems Agency

Air Force Research Laboratory, Space Vehicles Directorate

Air Force Research Laboratory, Information Systems Directorate

Airspace Command, Control, Intelligence, Surveillance and Reconnaissance Center

Joint Intelligence Center, Pacific Command

Lockheed-Martin Mission Systems

(This Page Intentionally Left Blank)

Initial Distribution

Headquarters Air Force

SAF/OS	Secretary of the Air Force
AF/CC	Chief of Staff
AF/CV	Vice Chief of Staff
AF/CVA	Assistant Vice Chief of Staff
AF/HO	Historian
AF/ST	Chief Scientist
AF/SG	Surgeon General
AF/SF	Security Forces
AF/TE	Test and Evaluation
AF/XI	Warfighting Integration

Assistant Secretary of the Air Force

SAF/AQ	Assistant Secretary for Acquisition
SAF/AQ	Military Director, USAF Scientific Advisory Board
SAF/AQI	Information Dominance
SAF/AQL	Special Programs
SAF/AQP	Global Power
SAF/AQQ	Global Reach
SAF/AQR	Science, Technology and Engineering
SAF/AQS	Space and Nuclear Deterrence
SAF/AQX	Management Policy and Program Integration
SAF/MI	Assistant Secretary (Manpower, Reserve Affairs, Installations & Environment)
SAF/SN	Assistant Secretary (Space)
SAF/SX	Deputy Assistant Secretary (Space Plans and Policy)

Deputy Chief of Staff, Air and Space Operations

AF/XO	DCS, Air and Space Operations
AF/XOC	Command and Control
AF/XOI	Intelligence, Surveillance, and Reconnaissance
AF/XOJ	Joint Matters
AF/XOO	Operations and Training
AF/XOR	Operational Requirements

Deputy Chief of Staff, Installations and Logistics

AF/IL	DCS, Installations and Logistics
AF/ILX	Plans and Integration

Deputy Chief of Staff, Plans and Programs

AF/XP	DCS, Plans and Programs
AF/XPI	Information and Systems
AF/XPM	Manpower, Organization and Quality
AF/XPP	Programs
AF/XPX	Strategic Planning
AF/XPY	Analysis

Initial Distribution (continued)

Deputy Chief of Staff, Personnel

AF/DP	DCS, Personnel

Office of the Secretary of Defense

USD (A&T)	Under Secretary for Acquisition and Technology
USD (A&T)/DSB	Defense Science Board
DARPA	Defense Advanced Research Projects Agency
DIA	Defense Intelligence Agency
DISA	Defense Information Systems Agency
BMDO	Ballistic Missile Defense Organization

Other Air Force Organizations

AC2ISRC	Aerospace Command, Control, Intelligence, Surveillance, and Reconnaissance Center
ACC	Air Combat Command
– CC	- Commander, Air Combat Command
– 366th Wing	- 366th Wing at Mountain Home Air Force Base
AETC	Air Education and Training Command
– AU	- Air University
AFMC	Air Force Materiel Command
– CC	- Commander, Air Force Materiel Command
– EN	- Directorate of Engineering and Technical Management
– AFRL	- Air Force Research Laboratory
– SMC	- Space and Missile Systems Center
– ESC	- Electronic Systems Center
– ASC	- Aeronautics Systems Center
– HSC	- Human Systems Center
– AFOSR	- Air Force Office of Scientific Research
AFOTEC	Air Force Operational Test and Evaluation Center
AFSAA	Air Force Studies and Analyses Agency
AFSOC	Air Force Special Operations Command
AFSPC	Air Force Space Command
AIA	Air Intelligence Agency
AMC	Air Mobility Command
NAIC	National Air Intelligence Center
NGB/CF	National Guard Bureau
PACAF	Pacific Air Forces
USAFA	U.S. Air Force Academy
USAFE	U.S. Air Forces in Europe

U.S. Army

ASB	Army Science Board

U.S. Navy

NRAC	Naval Research Advisory Committee
Naval Studies Board	

Initial Distribution (continued)

U.S. Marine Corps

DC/S (A)	Deputy Chief of Staff for Aviation

Joint Staff

JCS	Office of the Vice Chairman
J2	Intelligence
J3	Operations
J4	Logistics
J5	Strategic Plans and Policies
J6	Command, Control, Communications, and Computer Systems
J7	Operational Plans and Interoperability
J8	Force Structure, Resources and Assessment

Other

Aerospace Corporation
ANSER
MITRE
RAND
Study Participants

(This Page Intentionally Left Blank)

<table>
<tr><td colspan="4">REPORT DOCUMENTATION PAGE</td><td>Form Approved
OMB No. 0704-0188</td></tr>
<tr><td colspan="5">Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and manipulating the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget Paperwork Reduction Project (0704-0188), Washington, DC 20503</td></tr>
<tr><td>1. AGENCY USE ONLY (Leave Blank)</td><td colspan="2">2. REPORT DATE
June 2002</td><td colspan="2">3. REPORT TYPE AND DATES COVERED
Final, January 2000 - November 2000</td></tr>
<tr><td colspan="4">4. TITLE AND SUBTITLE
Database Migration for Command and Control</td><td rowspan="2">5. FUNDING NUMBERS</td></tr>
<tr><td colspan="4">6. AUTHOR(S)
Dr. James Hendler, Dr. Bhavani Thuraisingham, Maj Gen (r) Eric Nelson, Prof. Gio Wiederhold, Mr. Scott Saunders, Mr. Scott Fouse</td></tr>
<tr><td colspan="4">7. PERFORMING ORGANIZATION NAMES(S) AND ADDRESS(ES)
AF/SB
Pentagon
Washington, DC 20330-1180</td><td>8. PERFORMING ORGANIZATION REPORT NUMBER
SAB-TR-01-03</td></tr>
<tr><td colspan="4">9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
SAF/OS
AF/CC
Pentagon
Washington, DC 20330-1670</td><td>10. SPONSORING/MONITORING AGENCY REPORT NUMBER</td></tr>
<tr><td colspan="5">11. SUPPLEMENTARY NOTES</td></tr>
<tr><td colspan="4">12a. DISTRIBUTION/AVAILABILITY STATEMENT
Cleared for Open Publication</td><td>12b. DISTRIBUTION CODE</td></tr>
<tr><td colspan="5">ABSTRACT (Maximum 200 Words)
A report on the technologies, techniques and procedures required to manage effectively the transition of Air Force databases from one generation of information technology to the next.</td></tr>
<tr><td colspan="4" rowspan="2">14. SUBJECT TERMS
Science and Engineering Planning, Science and Technology, Leveraging Partnerships, Database Technology, Development Planning, Database Migration,</td><td>15. NUMBER OF PAGES
156</td></tr>
<tr><td>16. PRICE CODE</td></tr>
<tr><td>17. SECURITY CLASSIFICATION OF REPORT
Unclassified</td><td>18. SECURITY CLASSIFICATION OF THIS PAGE
Unclassified</td><td colspan="2">19. SECURITY CLASSIFICATION OF ABSTRACT
Unclassified</td><td>20. LIMITATION OF ABSTRACT
None</td></tr>
</table>

(This Page Intentionally Left Blank)